Junos® OS

Interfaces Fundamentals for Routing Devices

Modified: 2018-06-25

Copyright © 2018, Juniper Networks, Inc.

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

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. and/or its affiliates in the United States and other countries. All other trademarks may be property of their respective owners.

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

Junos® OS Interfaces Fundamentals for Routing Devices
Copyright © 2018 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 https://www.juniper.net/support/eula/. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

Copyright © 2018, Juniper Networks, Inc.

 Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

Part 5 Configuration Statements and Operational Commands

Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

 Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

 Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

 Copyright © 2018, Juniper Networks, Inc.

 Copyright © 2018, Juniper Networks, Inc.

 Copyright © 2018, Juniper Networks, Inc.

 Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

 Copyright © 2018, Juniper Networks, Inc.

 Copyright © 2018, Juniper Networks, Inc.

List of Figures

Copyright © 2018, Juniper Networks, Inc.

 Copyright © 2018, Juniper Networks, Inc.

List of Tables

Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

 Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

About the Documentation

Documentation and Release Notes

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

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 https://www.juniper.net/books.

Supported Platforms

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

- ACX Series
- M Series
- MX Series
- T Series
- PTX Series

Using the Examples in This Manual

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

Copyright © 2018, Juniper Networks, Inc.

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

 Copyright © 2018, Juniper Networks, Inc.

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

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

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

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

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

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

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

Documentation Conventions

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

Table 1: Notice Icons

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

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

Table 2: Text and Syntax Conventions

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

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

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

Copyright © 2018, Juniper Networks, Inc.

Table 2: Text and Syntax Conventions (continued)

Convention	Description	Examples
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	• In the Logical Interfaces box, select **All Interfaces**. • To cancel the configuration, click **Cancel**.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select **Protocols>Ospf**.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback rating system—On any page of the Juniper Networks TechLibrary site at https://www.juniper.net/documentation/index.html, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at https://www.juniper.net/documentation/feedback/.
- E-mail—Send your comments to **techpubs-comments@juniper.net**. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

Copyright © 2018, Juniper Networks, Inc.

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

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

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 https://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 https://www.juniper.net/support/requesting-support.html.

 Copyright © 2018, Juniper Networks, Inc.

PART 1

Router Interfaces

Copyright © 2018, Juniper Networks, Inc.

CHAPTER 1

Router Interfaces Overview

Router Interfaces Overview

Routers typically contain several different types of interfaces suited to various functions. For the interfaces on a router to function, you must configure them. Specify the interface location (that is, the slot where the Flexible PIC Concentrator [FPC], Dense Port Concentrator [DPC], or Modular Port Concentrator [MPC] is installed. You must also specify the location of the Physical Interface Card [PIC] or Modular Interface Card [MIC] , and the interface type, for example, SONET/SDH, Asynchronous Transfer Mode [ATM], or Ethernet). Finally, you must specify the encapsulation type and any interface-specific properties that may apply.

You can configure interfaces that are currently present in the router, as well as interfaces that are not currently present but that are expected to be added in the future. Junos OS detects the interface once the hardware has been installed and applies the pre-set configuration to it.

To see which interfaces are currently installed in the router, issue the **show interfaces terse** operational mode command. If an interface is listed in the output, it is physically installed in the router. If an interface is not listed in the output, it is not installed in the router.

For information about which interfaces are supported on your router, see your router's *Interface Module Reference*.

You can configure Junos OS class-of-service (CoS) properties to provide a variety of classes of service for different applications, including multiple forwarding classes for managing packet transmission, congestion management, and CoS-based forwarding. For more information about configuring CoS properties, see the *Class of Service Feature Guide for Routing Devices and EX9200 Switches*.

Related Documentation

- *Interfaces Fundamentals for Routing Devices*

Types of Interfaces Overview

Interfaces can be permanent or transient, and are used for networking or services:

- Permanent interfaces—Interfaces that are always present in the router.
- Transient interfaces—Interfaces that can be inserted into or removed from the router depending on your network configuration needs.
- Networking interfaces—Interfaces, such as Ethernet or SONET/SDH interfaces, that primarily provide traffic connectivity.
- Services interfaces—Interfaces that provide specific capabilities for manipulating traffic before it is delivered to its destination.
- Container interfaces—Interfaces that support automatic protection switching (APS) on physical SONET links using a virtual container infrastructure.

Copyright © 2018, Juniper Networks, Inc.

Junos OS internally generates nonconfigurable interfaces which are described in *Interfaces Command Reference* and *Services Interfaces*.

Related Documentation

- Understanding Permanent Interfaces on page 5
- Understanding Transient Interfaces on page 5
- Understanding Services Interfaces on page 7 and *Junos OS Services Interfaces Library for Routing Devices*
- Understanding Container Interfaces on page 8
- See also the following sections regarding specific networking interface technologies used in your routers:
  - *ATM Interfaces Overview*
  - *Channelized Interfaces Overview*
  - Circuit Emulation Interfaces: *Understanding Mobile Backhaul*
  - *E1 Interfaces Overview* and *E3 Interfaces Overview*
  - *Ethernet Interfaces Overview*
  - *Frame Relay Overview*
  - *SONET/SDH Interfaces Overview*
  - *T1 Interfaces Overview* and *T3 Interfaces Overview*

Understanding Permanent Interfaces

Permanent interfaces in the router consist of management Ethernet interfaces and internal Ethernet interfaces, which are described separately in the following topics:

- Understanding Management Ethernet Interfaces on page 10
- Understanding Internal Ethernet Interfaces on page 15

Understanding Transient Interfaces

The M Series, MX Series, and T Series routers contain slots for installing Flexible PIC Concentrator [FPC] or Dense Port Concentrator [DPC] (for MX Series routers) or Modular Port Concentrator [MPC] (for MX Series routers). Physical Interface Card [PIC] can be installed in FPCs. Modular Interface Card [MIC] can be inserted into MPCs.

The number of PICs that can be installed varies by router and type of FPC. The PICs provide the actual physical interfaces to the network. The MX Series routers contain slots for installing either DPC boards that provide the physical interfaces to the network or for installing FPCs in which PICs can be installed.

You can insert any DPC or FPC into any slot that supports them in the appropriate router. Typically, you can place any combination of PICs, compatible with your router, in any location on an FPC. (You are limited by the total FPC bandwidth, and by the fact that

Copyright © 2018, Juniper Networks, Inc.

some PICs physically require two or four of the PIC locations on the FPC. In some cases, power limitations or microcode limitations may also apply.) To determine DPC and PIC compatibility, see the see your router's *Interface Module Reference*.

You can insert MPC into any slot that supports them in the appropriate router. You can install up to two MICs of different media types in the same MPC as long as the MPC supports those MICs.

These physical interfaces are transient interfaces of the router. They are referred to as transient because you can hot-swap a DPC or FPC or MPC and its PICs or MICs at any time.

You must configure each transient interface based on the slot in which the FPC or DPC or MPC is installed, the location in which the PIC or MIC is installed, and for multiple port PICs or MICs , the port to which you are connecting.

You can configure the interfaces on PICs or MICs that are already installed in the router as well as interfaces on PICs or MICs that you plan to install later. The Junos OS detects which interfaces are actually present, so when the software activates its configuration, it activates only the present interfaces and retains the configuration information for the interfaces that are not present. When the Junos OS detects that an FPC containing PICs or MPC containing MICs has been inserted into the router, the software activates the configuration for those interfaces.

Related Documentation

- Types of Interfaces Overview on page 4
- Understanding Permanent Interfaces on page 5
- Understanding Management Ethernet Interfaces on page 10
- Understanding Internal Ethernet Interfaces on page 15
- Supported Routing Engines by Router on page 17
- Understanding Services Interfaces on page 7
- Understanding Container Interfaces on page 8
- Interface Encapsulations Overview on page 46
- Interface Descriptors Overview on page 58
- Interface Naming Overview on page 32
- Displaying Interface Configurations Overview on page 61

 Copyright © 2018, Juniper Networks, Inc.

Understanding Services Interfaces

Services interfaces enable you to incrementally add services to your network. The Junos OS supports the following services PICs:

- Adaptive Services (AS) PICs—Allow you to provide multiple services on a single PIC by configuring a set of services and applications. The AS PICs offer a special range of services you configure in one or more service sets.
- ES PIC—Provides a security suite for the IP version 4 (IPv4) and IP version 6 (IPv6) network layers. The suite provides functionality such as authentication of origin, data integrity, confidentiality, replay protection, and nonrepudiation of source. It also defines mechanisms for key generation and exchange, management of security associations, and support for digital certificates.
- Monitoring Services PICs—Enable you to monitor traffic flow and export the monitored traffic. Monitoring traffic allows you to gather and export detailed information about IPv4 traffic flows between source and destination nodes in your network; sample all incoming IPv4 traffic on the monitoring interface and present the data in cflowd record format; perform discard accounting on an incoming traffic flow; encrypt or tunnel outgoing cflowd records, intercepted IPv4 traffic, or both; and direct filtered traffic to different packet analyzers and present the data in its original format. On a Monitoring Services II PIC, you can configure either monitoring interfaces or collector interfaces. A collector interface allows you to combine multiple cflowd records into a compressed ASCII data file and export the file to an FTP server.
- Multilink Services, MultiServices, Link Services, and Voice Services PICs—Enable you to split, recombine, and sequence datagrams across multiple logical data links. The goal of multilink operation is to coordinate multiple independent links between a fixed pair of systems, providing a virtual link with greater bandwidth than any of the members.
- Tunnel Services PIC—By encapsulating arbitrary packets inside a transport protocol, tunneling provides a private, secure path through an otherwise public network. Tunnels connect discontinuous subnetworks and enable encryption interfaces, virtual private networks (VPNs), and Multiprotocol Label Switching (MPLS).
- On M Series and T Series routers, logical tunnel interfaces allow you to connect logical systems, virtual routers, or VPN instances. For more information about VPNs, see the *Junos OS VPNs Library for Routing Devices*. For more information about configuring tunnels, see the *Junos OS Services Interfaces Library for Routing Devices*.

Related Documentation

- Types of Interfaces Overview on page 4

Understanding Container Interfaces

Container interfaces provide the following features:

- Automatic protection switching (APS) on SONET/SDH and ATM links are supported using the container infrastructure.
- Container physical interfaces and logical interfaces remain up on switchover.
- APS parameters are auto-copied from the container interface to the member links.

NOTE: Paired groups and true unidirectional APS are not currently supported.

For more information on SONET/SDH configuration, see *Configuring Container Interfaces for APS on SONET Links*.

Container interfaces features are described in the following sections:

Understanding Traditional APS Concept

Traditional APS is configured on two independent physical SONET/SDH interfaces: one configured as the working circuit and the other as the protect circuit (see Figure 1 on page 8). The circuit, named Circuit X in the figure, is the link between the two SONET interfaces.

Figure 1: APS Interface

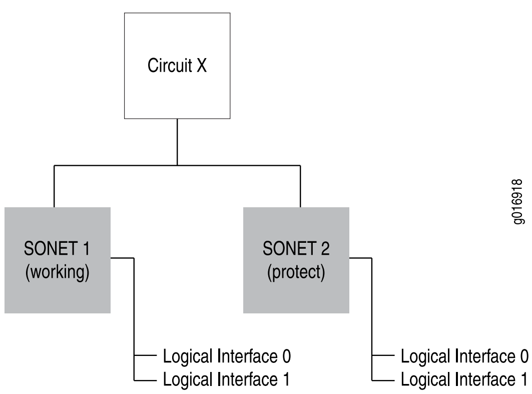

Traditional APS uses routing protocols that run on each individual SONET/SDH interface (since circuit is an abstract construct, instead of being an actual interface). When the working link goes down, the APS infrastructure brings up the protect link and its underlying logical interfaces, and brings down the working link and its underlying logical interfaces,

 Copyright © 2018, Juniper Networks, Inc.

causing the routing protocols to reconverge. This consumes time and leads to traffic loss even though the APS infrastructure has performed the switch quickly.

Container Interfaces Concept

To solve this problem, the Junos OS provides a soft interface construct called a container interface (see Figure 2 on page 9).

Figure 2: Container Interface

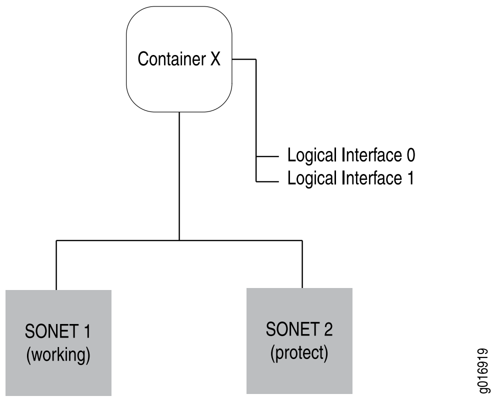

The container interface allows routing protocols to run on the logical interfaces associated with a virtual *container interface* instead of on the physical SONET/SDH and ATM interfaces. When APS switches the underlying physical link based on a fault condition, the container interface remains up, and the logical interface on the container interface does not flap. The routing protocols remain unaware of the APS switching.

APS Support for Container-Based Interfaces

With the container interface, APS is configured on the container interface itself. Individual member SONET/SDH and ATM links are either marked as primary (corresponding to the working circuit) or standby (corresponding to the protect circuit) in the configuration. No circuit or group name is specified in the container interface model; physical SONET/SDH and ATM links are put in an APS group by linking them to a single container interface. APS parameters are specified at the container interface level, and are propagated to the individual SONET/SDH and ATM links by the APS daemon.

Autocopy of APS Parameters

Typical applications require copying APS parameters from the working circuit to the protect circuit, since most of the parameters must be the same for both circuits. This is automatically done in the container interface. APS parameters are specified only once under the container physical interface configuration, and are internally copied over to the individual physical SONET/SDH and ATM links.

Related Documentation

- *Configuring Container Interfaces for APS on SONET Links*
- *Displaying APS Using a Container Interface with ATM Encapsulation*

Copyright © 2018, Juniper Networks, Inc.

Understanding Management Ethernet Interfaces

Management interfaces are the primary interfaces for accessing the device remotely. Typically, a management interface is not connected to the in-band network, but is connected instead to the device's internal network. Through a management interface you can access the device over the network using utilities such as **ssh** and **telnet** and configure the device from anywhere, regardless of its physical location. SNMP can use the management interface to gather statistics from the device.

A management interface lets authorized users and management systems connect to the device over the network. Some Juniper Networks devices have a dedicated management port on the front panel. For other types of platforms, you can configure a management interface on one of the network interfaces. This interface can be dedicated to management or shared with other traffic. Before users can access the management interface, you must configure it. Information required to set up the management interface includes its IP address and prefix. In many types of Junos OS devices (or recommended configurations), it is not possible to route traffic between the management interface and the other ports. Therefore, you should select an IP address in a separate (logical) network, with a separate prefix (netmask).

For devices with dedicated management ports, Junos OS automatically configures the router's management Ethernet interface, as either em0 or fxp0. You can use the **show interfaces terse | match fxp0** or **show interfaces terse | match em0** command to display management interface information.

To use the management Ethernet interface as a management port, you must configure its logical port, em0.0 or fxp0.0, with a valid IP address.

For some SRX Series Services Gateways and J Series Services Routers, you can use any of the built-in Ethernet ports as a management interface. (Platform support depends on the Junos OS release in your installation.) To use a built-in interface as a management Ethernet interface, configure it with a valid IP address. To manually configure J-Web access, include the **interface *interface-name*** statement at the **[edit system services web-management http]** hierarchy level.

For PTX Series Packet Transport Routers, the Junos OS automatically creates the router's management Ethernet interface, **em0**. To use **em0** as an out-of-band management port, you must configure its logical port (for example, **em0.0**) with a valid IP address.

Internal Ethernet interfaces are automatically created to connect the Routing Engines to the Packet Forwarding Engines in the FPCs.

When you enter the **show interfaces** command on a PTX Series Packet Transport Router, the management Ethernet interface and internal Ethernet interfaces (and logical interfaces) are displayed:

```
user@host> show interfaces ?
...
em0
  em0.0
  ixgbe0
```

 Copyright © 2018, Juniper Networks, Inc.

```
  ixgbe0.0
  ixgbe1
  ixgbe1.0
...
```

NOTE: *Routing Engine upgrade considerations*—When upgrading to a Routing Engine that supports em0 from a Routing Engine that supports fxp0, you must convert existing management Ethernet interface references in the router configuration files from fxp0, fxp1, or fxp2 interfaces to em0 interfaces. Whether you use an automated script or edit the configuration files manually, you must revise any command lines that reference the fxp0 management Ethernet interface by replacing "fxp0" with "em0."

Reusing scripts for standalone T1600 routers on T1600 routers in a routing matrix—Automated scripts that you have developed for standalone T1600 routers (T1600 routers that are not in a routing matrix) might contain references to the fxp0 management Ethernet interface. Before reusing the scripts on T1600 routers in a routing matrix, edit the command lines that reference the fxp0 management Ethernet interface so that the commands reference the em0 management Ethernet interface instead.

Restricted load-sharing next hops with fxp0—On M Series Multiservices Edge Routers and T Series Core Routers running Junos OS later than Release 7.0R2.7 or Release 7.1R2.2, the fxp0 interface does not support load-sharing next hops. This restriction only affects fxp0 routes.

CoS not supported on fxp0—The fxp0 interface does not support class of service (CoS).

The Routing Engines in the PTX Series Packet Transport Routers do not support the management Ethernet interface **fxp0**, or the internal Ethernet interfaces **fxp1** or **fxp2**.

Related Documentation

- Supported Routing Engines by Router on page 17

Understanding Interfaces on ACX Series Universal Metro Routers

The ACX Series routers support time-division multiplexing (TDM) T1 and E1 interfaces and Ethernet (1 GbE copper, 1GbE, 10 GbE, and 40 GbE fiber) interfaces to support both the legacy and evolution needs of the mobile network. Support for Power over Ethernet (PoE+) at 65 watts per port mitigates the need for additional electrical cabling for microwaves or other access interfaces.

The ACX Series routers support the following:

- TDM T1 and E1 ports:
  - The ACX1000 router contains eight T1 or E1 ports.

Copyright © 2018, Juniper Networks, Inc.

- The ACX2000 router contains 16 T1 or E1 ports.
- Inverse Multiplexing for ATM (IMA)

NOTE: **ACX5048 and ACX5096 routers do not support T1 or E1 ports and Inverse Multiplexing for ATM (IMA).**

- Gigabit Ethernet ports:
  - The ACX1000 router contains eight Gigabit Ethernet ports. The ACX1000 router also supports either four RJ45 (Cu) ports or installation of four Gigabit Ethernet small form-factor pluggable (SFP) transceivers.
  - The ACX2000 router contains 16 Gigabit Ethernet ports and two PoE ports. The ACX2000 router also supports installation of two Gigabit Ethernet SFP transceivers and two 10-Gigabit Ethernet SFP+ transceivers.
  - The ACX5448 router is a 10-Gigabit Ethernet enhanced small form-factor pluggable (SFP+) top-of-rack router with 48 SFP+ ports, and four 100-Gigabit Ethernet QSFP28 ports. Each SFP+ port can operate as a native 10-Gigabit Ethernet port, or as a 1-Gigabit Ethernet port when 1-Gigabit optics are inserted. The 48 ports on ACX5448 router can be configured as 1GE or 10GE modes and these ports are represented by **xe** interface type. The PIC 1 of FPC 0 has 4x100GE ports, where each port can be channelized as 1x100GE, or 1x40GE, or 4x25GE modes and these ports are represented by **et** interface type. By default, the port speed in PIC 1 is 100GE.

    NOTE: **The ACX5448 router do not support Pseudowire Services interface.**

NOTE: **40GbE is supported only on ACX5048 and ACX5096 routers.**

T1 and E1 Time-Division Multiplexing (TDM) Interfaces

On the ACX Series routers, existing Junos OS TDM features are supported without changes to statements or functionality. The following key TDM features for T1 (**ct1**) interfaces and E1 (**ce1**) interfaces are supported:

- T1 and E1 channelization
- T1 and E1 encapsulation
- Alarms, defects, and statistics
- External and internal loopback
- TDM class of service (CoS)

 Copyright © 2018, Juniper Networks, Inc.

T1 and E1 mode selection is at the PIC level. To set the T1 or E1 mode at the PIC level, include the **framing** statement with the **t1** or **e1** option at the [**chassis fpc *slot-number* pic *slot-number***] hierarchy level. All ports can be T1 or E1. Mixing T1s and E1s is not supported.

T1 or E1 BITS Interface (ACX2000)

The ACX2000 router has a T1 or E1 building-integrated timing supply (BITS) interface that you can connect to an external clock. After you connect the interface to the external clock, you can configure the BITS interface so that the BITS interface becomes a candidate source for chassis synchronization to the external clock. The frequency of the BITS interface depends on the Synchronous Ethernet equipment slave clock (EEC) selected with the **network-option** statement at the [**edit chassis synchronization**] hierarchy level.

NOTE: The ACX1000 router does not support the BITS interface.

Inverse Multiplexing for ATM (IMA)

Defined by the ATM Forum, IMA specification version 1.1 is a standardized technology used to transport ATM traffic over a bundle of T1 and E1 interfaces, also known as an IMA group. Up to eight links per bundle and 16 bundles per PIC are supported. The following key IMA features are supported:

- IMA Layer 2 encapsulation
- ATM CoS
- ATM policing and shaping
- Denied packets counter in the output for the **show interfaces at-*fpc/pic/port* extensive** command

Gigabit Ethernet interfaces

On the ACX Series routers, existing Junos OS Ethernet features are supported without changes to statements or functionality. The following key features are supported:

- Media type specification (ACX1000 router with Gigabit Ethernet SFP and RJ45 interfaces)
- Autonegotiation for RJ45 Gigabit Ethernet interfaces
- Event handling of SFP insertion and removal
- Explicit disabling of the physical interface
- Flow control

  NOTE: The ACX Series router does not support flow control based on PAUSE frames.

- Loopback

Copyright © 2018, Juniper Networks, Inc.

- Loss of signal (LOS) alarm
- Media access control (MAC) layer features
- Maximum transmission unit (MTU)
- Remote fault notification for 10-Gigabit Ethernet interfaces
- Statistics collection and handling
- Power over Ethernet (PoE) (ACX2000 router)
- High power mode

The Gigabit Ethernet ports on the router have the capacity to work as a 1 or 10-Gigabit Ethernet interface, depending on the type of small form-factor pluggable (SFP) transceiver inserted. When you insert an SFP+ transceiver, the interface works at the 10-Gigabit speed. When you insert an SFP transceiver, the interface works at the 1-Gigabit speed. Configuration is not required because the speed is determined automatically based on the type of inserted SFP transceiver. The dual-speed interface is automatically created with the **xe** prefix, for example, **xe-4/0/0**.

The same configuration statements are used for both speeds and CoS parameters are scaled as a percentage of the port speed. To configure a dual-speed Gigabit Ethernet interface, include the **interface xe-*fpc/pic/port*** statement at the [**edit interfaces**] hierarchy level. To display the interface speed and other details, issue the **show interfaces** command.

NOTE: You need to use industrial grade of SFP below 0dC for ACX 1100 and ACX 2100 boards.

Related Documentation

- *Understanding Encapsulation on an Interface*
- *Configuring Inverse Multiplexing for ATM (IMA) on ACX Series*
- Interface Names for ACX Series Universal Metro Routers on page 59

TX Matrix Plus and T1600 Router (Routing Matrix) Management Ethernet Interfaces

For TX Matrix Plus Routers and for T1600 Core Routers with RE-C1800 configured in a routing matrix, the Junos OS automatically creates the router's management Ethernet interface, **em0**. To use **em0** as a management port, you must configure its logical port, **em0.0**, with a valid IP address.

When you enter the **show interfaces** command on a TX Matrix Plus router, the management Ethernet interfaces (and logical interfaces) are displayed:

```
user@host> show interfaces ?
...
em0
em0.0
...
```

 Copyright © 2018, Juniper Networks, Inc.

NOTE: **The Routing Engines in the TX Matrix Plus router and in the T1600 routers with RE-C1800 configured in a routing matrix do not support the management Ethernet interface fxp0, or the internal Ethernet interfaces fxp1 or fxp2.**

Related Documentation

- Understanding Internal Ethernet Interfaces on page 15
- T1600 Routers (Routing Matrix) Internal Ethernet Interfaces on page 17
- *Displaying Internal Ethernet Interfaces for a Routing Matrix with a TX Matrix Plus Router*
- **show interfaces (M Series, MX Series, T Series Routers, and PTX Series Management and Internal Ethernet) on page 1422**

Understanding Internal Ethernet Interfaces

Within a router or packet transport router, internal Ethernet interfaces provide communication between the Routing Engine and the Packet Forwarding Engines. The Junos OS automatically configures internal Ethernet interfaces when the Junos OS boots. The Junos OS boots the packet-forwarding component hardware. When these components are running, the Control Board uses the internal Ethernet interface to transmit hardware status information to the Routing Engine. Information transmitted includes the internal router temperature, the condition of the fans, whether an FPC has been removed or inserted, and information from the LCD on thecraft interface.

To determine the supported internal Ethernet interfaces for your router, see "Supported Routing Engines by Router" on page 17.

NOTE: **Do not modify or remove the configuration for the internal Ethernet interface that the Junos OS automatically configures. If you do, the router or packet transport routerwill stop functioning.**

- M Series, and MX Series routers and T Series routers—The Junos OS creates the internal Ethernet interface. The internal Ethernet interface connects the Routing Engine **re0** to the Packet Forwarding Engines.

  If the router has redundant Routing Engines, another internal Ethernet interface is created on each Routing Engine (**re0** and **re1**) in order to support fault tolerance, two physical links between **re0** and **re1** connect the independent control planes. If one of the links fails, both Routing Engines can use the other link for IP communication.

- TX Matrix Plus routers—On a TX Matrix Plus router, the Routing Engine and Control Board function as a unit, or host subsystem. For each host subsystem in the router, the Junos OS automatically creates two internal Ethernet interfaces, **ixgbe0** and **ixgbe1**.

Copyright © 2018, Juniper Networks, Inc.

The ixgbe0 and ixgbe1 interfaces connect the TX Matrix Plus Routing Engine to the Routing Engines of every line-card chassis (LCC) configured in the routing matrix.

The TX Matrix Plus Routing Engine connects to a high-speed switch through a 10-Gbps link within the host subsystem. The switch provides a 1-Gbps link to each T1600 Routing Engine. The 1-Gbps links are provided through the UTP Category 5 Ethernet cable connections between the TXP-CBs and the LCC-CBs in the LCCs.

- The TX Matrix Plus Routing Engine connects to a high-speed switch in the local Control Board through a 10-Gbps link within the host subsystem.
- The Gigabit Ethernet switch connects the Control Board to the remote Routing Engines of every LCC configured in the routing matrix.

If a TX Matrix Plus router contains redundant host subsystems, the independent control planes are connected by two physical links between the two 10-Gigabit Ethernet ports on their respective Routing Engines.

- The primary link to the remote Routing Engine is at the **ixgbe0** interface; the 10-Gigabit Ethernet switch on the local Control Board also connects the Routing Engine to the 10-Gigabit Ethernet port accessed by the **ixgbe1** interface on the remote Routing Engine.
- The alternate link to the remote Routing Engine is the 10-Gigabit Ethernet port at the **ixgbe1** interface. This second port connects the Routing Engine to the 10-Gigabit Ethernet switch on the remote Control Board, which connects to the 10-Gigabit Ethernet port at the **ixgbe0** interface on the remote Routing Engine.

If one of the two links between the host subsystems fails, both Routing Engines can use the other link for IP communication.

- LCC in a routing matrix—On an LCC configured in a routing matrix, the Routing Engine and Control Board function as a unit, or host subsystem. For each host subsystem in the LCC, the Junos OS automatically creates two internal Ethernet interfaces, **bcm0** and **em1**, for the two Gigabit Ethernet ports on the Routing Engine.

  The **bcm0** interface connects the Routing Engine in each LCCto the Routing Engines of every other LCC configured in the routing matrix.

  - The Routing Engine connects to a Gigabit Ethernet switch on the local Control Board through a.
  - The switch connects the Control Board to the remote Routing Engines of every other LCC configured in the routing matrix.

  If an LCC in a routing matrix contains redundant host subsystems, the independent control planes are connected by two physical links between the Gigabit Ethernet ports on their respective Routing Engines.

  - The primary link to the remote Routing Engine is at the **bcm0** interface; the Gigabit Ethernet switch on the local Control Board also connects the Routing Engine to the Gigabit Ethernet port accessed by the **em1** interface on the remote Routing Engine.

 Copyright © 2018, Juniper Networks, Inc.

- The alternate link to the remote Routing Engine is at the **em1** interface. This second port connects the Routing Engine to the Gigabit Ethernet switch on the remote Control Board, which connects to the Gigabit Ethernet port at the **bcm0** interface on the remote Routing Engine.

If one of the two links between the host subsystems fails, both Routing Engines can use the other link for IP communication.

Each router also has two serial ports, labeled *console* and *auxiliary*, for connecting tty type terminals to the router using standard PC-type tty cables. Although these ports are not network interfaces, they do provide access to the router.

Related Documentation

- Understanding Permanent Interfaces on page 5
- Supported Routing Engines by Router on page 17
- TX Matrix Plus and T1600 Router (Routing Matrix) Management Ethernet Interfaces on page 14
- T1600 Routers (Routing Matrix) Internal Ethernet Interfaces on page 17
- *Displaying Internal Ethernet Interfaces for a Routing Matrix with a TX Matrix Plus Router*
- **show interfaces (M Series, MX Series, T Series Routers, and PTX Series Management and Internal Ethernet) on page 1422**

T1600 Routers (Routing Matrix) Internal Ethernet Interfaces

On a T1600 router configured in a routing matrix, the Routing Engine (RE-TXP-LCC) and Control Board (LCC-CB) function as a unit, or host subsystem. For each host subsystem in the router, the Junos OS automatically creates two internal Ethernet interfaces, **bcm0** and **em1**, for the two Gigabit Ethernet ports on the Routing Engine.

Related Documentation

- Understanding Internal Ethernet Interfaces on page 15
- *Displaying Internal Ethernet Interfaces for a Routing Matrix with a TX Matrix Plus Router*
- **show interfaces (M Series, MX Series, T Series Routers, and PTX Series Management and Internal Ethernet) on page 1422**

Supported Routing Engines by Router

The following tables list the Routing Engines that each router supports, the first supported release for the Routing Engine in the specified router, the management Ethernet interface, and the internal Ethernet interfaces for each Routing Engine.

- M7i Routing Engines on page 18
- M10i Routing Engines on page 19
- M40e Routing Engines on page 19
- M120 Routing Engines on page 19

M7i Routing Engines

Table 3 on page 18 lists the Routing Engines supported by the M7i router. The M7i router supports 32-bit Junos OS only.

Table 3: M7i Routing Engines

Model Number	Name in CLI Output	First Supported 32-bit Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
RE-400-768 (EOL details: TSB16445)	**RE-5.0**	9.0	fxp0	fxp1
RE-850-1536 (EOL details: TSB15553)	**RE-850**	7.2	fxp0	fxp1
RE-B-1800X1-4G	**RE-B-1800x1**	11.4R4 12.1R2	fxp0	em0

 Copyright © 2018, Juniper Networks, Inc.

M10i Routing Engines

Table 4 on page 19 lists the Routing Engines supported by the M10i router. The M10i router supports 32-bit Junos OS only.

Table 4: M10i Routing Engines

Model Number	Name in CLI Output	First Supported 32-bit Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
RE-400-768 (EOL details: TSB16445)	**RE-5.0**	9.0	fxp0	fxp1 fxp2
RE-850-1536 (EOL details: TSB15553)	**RE-850**	7.2	fxp0	fxp1 fxp2
RE-B-1800X1-4G	**RE-B-1800x1**	11.4R4 12.1R2	fxp0	em0

M40e Routing Engines

Table 5 on page 19 lists the Routing Engines supported by the M40e router.

Table 5: M40e Routing Engines

Model Number	Name in CLI Output	First Supported Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
RE-600-2048 (EOL details: TSB14373)	**RE-3.0** or **RE-3.0 (RE-600)**	5.3	fxp0	fxp1 fxp2
RE-A-1000-2048	**RE-A-1000**	8.1	fxp0	fxp1 fxp2

M120 Routing Engines

Table 6 on page 19 lists the Routing Engines supported by the M120 router.

Table 6: M120 Routing Engines

Model Number	Name in CLI Output	First Supported 32-bit Junos OS Release	First Supported 64-bit Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
RE-A-1000-2048	**RE-A-1000**	8.0R2	–	fxp0	fxp1 fxp2

Copyright © 2018, Juniper Networks, Inc.

Table 6: M120 Routing Engines (continued)

Model Number	Name in CLI Output	First Supported 32-bit Junos OS Release	First Supported 64-bit Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
RE-A-2000-4096	**RE-A-2000**	8.0R2	–	fxp0	em0 bcm0
RE-A-1800X2-8G	**RE-A-1800x2**	• 11.4R5 • 12.1R3	10.4	fxp0	fxp1 fxp2
RE-A-1800X2-16G	**RE-A-1800x2**	• 11.4R5 • 12.1R3	10.4	fxp0	fxp1 fxp2
RE-A-1800X4-16G	**RE-A-1800x4**	• 11.4R5 • 12.1R3	10.4	fxp0	em0 em1

M320 Routing Engines

Table 7 on page 20 lists the Routing Engines supported by the M320 router.

Table 7: M320 Routing Engines

Model Number	Name in CLI Output	First Supported 32-bit Junos OS Release	First Supported 64-bit Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
RE-1600-2048 (EOL details: TSB14374)	**RE-4.0**	6.2	–	fxp0	fxp1 fxp2
RE-A-2000-4096	**RE-A-2000**	8.1	–	fxp0	em0 bcm0
RE-A-1800X2-8G	**RE-A-1800x2**	• 11.4R5 • 12.1R3	10.4	fxp0	em0 bcm0
RE-A-1800X2-16G	**RE-A-1800x2**	• 11.4R5 • 12.1R3	10.4	fxp0	em0 bcm0
RE-A-1800X4-8G	**RE-A-1800X4**	• 11.4R5 • 12.1R3 • 12.2	10.4	fxp0	em0 em1

Copyright © 2018, Juniper Networks, Inc.

MX5, MX10, MX40, and MX80 Routing Engine

Table 8 on page 21 lists the Routing Engines supported by the MX5, MX10, MX40, and MX80 routers.

Table 8: MX5, MX10, MX40, and MX80 Routing Engine

Model Number	Name in CLI Output	First Supported 32-bit Junos OS Release	First Supported 64-bit Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
Built-in Routing Engine	**Routing Engine RE-MX80**	12.3	-	fxp0	em0 em1 NOTE: em1 is used to communicate with the MS-MIC when it is inserted.

MX104 Routing Engines

Table 9 on page 21 lists the Routing Engines supported by MX104 routers.

Table 9: MX104 Routing Engines

Model Number	Name in CLI Output	First Supported 32-bit Junos OS Release	First Supported 64-bit Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
RE-S-MX104	**Routing Engine**	13.2	–	fxp0	fxp1 fxp2

MX240 Routing Engines

Table 10 on page 21 lists the Routing Engines supported by MX240 routers.

Table 10: MX240 Supported Routing Engines

Model Number	Name in CLI Output	First Supported 32-bit Junos OS Release	First Supported 64-bit Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
RE-S-1300-2048 (EOL details: TSB16556	**RE-S-1300**	9.0	–	fxp0	fxp1 fxp2
RE-S-2000-4096 (EOL details: TSB16735	**RE-S-2000**	9.0	–	fxp0	fxp1 fxp2
RE-S-1800X2-8G (EOL details: TSB16556	**RE-S-1800x2**	• 11.4R5 • 12.1R3	10.4	fxp0	em0 em1

Copyright © 2018, Juniper Networks, Inc.

Table 10: MX240 Supported Routing Engines (continued)

Model Number	Name in CLI Output	First Supported 32-bit Junos OS Release	First Supported 64-bit Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
RE-S-1800x2-16G (EOL details: TSB16556	**RE-S-1800x2**	• 11.4R5 • 12.1R3	10.4	fxp0	em0 em1
RE-S-1800X4-8G	**RE-S-1800X4**	• 11.4R5 • 12.1R3	10.4	fxp0	em0 em1
RE-S-1800X4-16G	**RE-S-1800x4**	• 11.4R5 • 12.1R3	10.4	fxp0	em0 em1
RE-S-1800X4-32G-S	**RE-S-1800X4**	• 12.3R4 • 13.2R1	• 12.3R4 • 13.2R1	fxp0	em0, em1
RE-S-X6-64G	**RE-S-2X00x6**	–	15.1F4 16.1R1	fxp0	ixlv0, igb0

MX480 Routing Engines

Table 11 on page 22 lists the Routing Engines supported by MX480 routers.

Table 11: MX480 Supported Routing Engines

Model Number	Name in CLI Output	First Supported 32-bit Junos OS Release	First Supported 64-bit Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
RE-S-1300-2048 (EOL details: TSB16556	**RE-S-1300**	8.4	–	fxp0	fxp1 fxp2
RE-S-2000-4096 (EOL details: TSB16735	**RE-S-2000**	8.4	–	fxp0	fxp1 fxp2
RE-S-1800X2-8G (EOL details: TSB16556	**RE-S-1800x2**	• 11.4R5 • 12.1R3	10.4	fxp0	em0 em1
RE-S-1800X2-16G (EOL details: TSB16556	**RE-S-1800x2**	• 11.4R5 • 12.1R3	10.4	fxp0	em0 em1
RE-S-1800X4-8G	**RE-S-1800X4**	• 11.4R5 • 12.1R3	10.4	fxp0	em0 em1

Copyright © 2018, Juniper Networks, Inc.

Table 11: MX480 Supported Routing Engines (continued)

Model Number	Name in CLI Output	First Supported 32-bit Junos OS Release	First Supported 64-bit Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
RE-S-1800X4-16G	**RE-S-1800x4**	• 11.4R5 • 12.1R3	10.4	fxp0	em0 em1
RE-S-1800X4-32G-S	**RE-S-1800X4**	• 12.3R4 • 13.2R1	• 12.3R4 • 13.2R1	fxp0	em0 em1
RE-S-X6-64G	**RE-S-2X00x6**	–	15.1F4 16.1R1	fxp0	ixlv0, igb0

MX960 Routing Engines

Table 12 on page 23 lists the Routing Engines supported by MX960 routers.

Table 12: MX960 Supported Routing Engines

Model Number	Name in CLI Output	First Supported 32-bit Junos OS Release	First Supported 64-bit Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
RE-S-1300-2048 (EOL details: TSB16556	**RE-S-1300**	8.2	–	fxp0	fxp1 fxp2
RE-S-2000-4096 (EOL details: TSB16735	**RE-S-2000**	8.2	–	fxp0	fxp1 fxp2
RE-S-1800X2-8G (EOL details: TSB16556	**RE-S-1800x2**	• 11.4R5 • 12.1R3	10.4	fxp0	em0 em1
RE-S-1800X2-16G (EOL details: TSB16556	**RE-S-1800x2**	• 11.4R5 • 12.1R3	10.4	fxp0	em0 em1
RE-S-1800X4-8G	**RE-S-1800x4**	• 11.4R5 • 12.1R3	10.4	fxp0	em0 em1
RE-S-1800X4-16G	**RE-S-1800x4**	• 11.4R5 • 12.1R3	10.4	fxp0	em0 em1
RE-S-1800X4-32G-S	**RE-S-1800x4**	• 12.3R4 • 13.2R1	• 12.3R4 • 13.2R1	fxp0	em0 em1

Copyright © 2018, Juniper Networks, Inc.

Table 12: MX960 Supported Routing Engines (continued)

Model Number	Name in CLI Output	First Supported 32-bit Junos OS Release	First Supported 64-bit Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
RE-S-X6-64G	**RE-S-2X00x6**	–	15.1F4 16.1R1	fxp0	ixlv0, igb0
RE-S-X6-64G (For MX960-VC) added for PR 1316954	**RE-S-2X00x6**	–	17.1R1	fxp0	ixlv0, igb0
RE-S-X6-128G	**RE-S-2X00x6-128**	–	18.1R1	fxp0	ixlv0, igb0 em0

MX2008 Routing Engines

Table 13 on page 24 lists the Routing Engines supported by MX2008 routers.

Table 13: MX2008 Supported Routing Engines

Model Number	Name in CLI Output	First Supported 64-bit Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
REMX2008-X8-64G	**RE-MX2008-X8-64G**	15.1F7	fxp0	ixlv0 ixlv1
REMX2008-X8-64G-LT	RE-MX2008-X8-64G-LT	17.2R1	fxp0	ixlv0 ixlv1
REMX2008-X8-128G	RE-MX2008-X8-128G	18.2R1	fxp0	ixlv0 ixlv1

MX2010 Routing Engines

Table 14 on page 24 lists the Routing Engines supported by MX2010 routers.

Table 14: MX2010 Supported Routing Engines

Model Number	Name in CLI Output	First Supported 64-bit Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
RE-MX2000-1800X4	**RE-S-1800x4**	12.3R2	fxp0	em0 em1

Copyright © 2018, Juniper Networks, Inc.

Table 14: MX2010 Supported Routing Engines (continued)

Model Number	Name in CLI Output	First Supported 64-bit Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
REMX2K-1800-32G-S	**RE-S-1800x4**	• 12.3R4 • 13.2R1	fxp0	em0 em1
REMX2K-X8-64G	**RE-S-2X00x8**	• 15.1F5-S1 • 16.1R2 • 16.2R1	fxp0	ixlv0 ixlv1 em0
REMX2008-X8-64G-LT	**RE-S-2X00x8-LT**	17.2R1	fxp0	ixlv0 ixlv1 em0
REMX2K-X8-128G	**RE-MX200X8-128G**	18.1R1	fxp0	ixlv0 ixlv1 em0

MX2020 Supported Routing Engines

Table 15 on page 25 lists the Routing Engines supported by MX2020 routers.

Table 15: MX2020 Supported Routing Engines

Model Number	Name in CLI Output	First Supported 64-bit Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
RE-MX2000-1800X4	**RE-S-1800x4**	12.3R2	fxp0	em0 em1
REMX2K-1800-32G-S	**RE-S-1800x4**	• 12.3R4 • 13.2R1	fxp0	em0 em1
REMX2K-X8-64G	**RE-S-2X00x8**	• 15.1F5-S1 • 16.1R2 • 16.2R1	fxp0	ixlv0 ixlv1 em0
REMX2008-X8-64G-LT	**RE-S-2X00x8-LT**	17.2R1	fxp0	ixlv0 ixlv1 em0

Copyright © 2018, Juniper Networks, Inc.

Table 15: MX2020 Supported Routing Engines (continued)

Model Number	Name in CLI Output	First Supported 64-bit Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
REMX2K-X8-128G	**RE-MX200X8-128G**	18.1R1	fxp0	ixlv0 ixlv1 em0

MX10003 Routing Engines

Table 16 on page 26 lists the Routing Engines supported by MX10003 routers.

Table 16: MX10003 Supported Routing Engines

Model Number	Name in CLI Output	First Supported 64-bit Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
JNP10003-RE1	**RE-S-2X00x6**	17.3R1	fxp0	ixlv0 ixlv1
JNP10003-RE1-LT		18.1R1	fxp0	ixlv0 ixlv1

MX10008 Routing Engines

Table 17 on page 26 lists the Routing Engines supported on the MX10008 router.

Table 17: MX10008 Routing Engines

Model Number	Name in CLI Output	First Supported Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
JNP10K-RE1	**JNP10K-RE1 RE X10**	18.2R1	em0	bme0 bme1

PTX1000 Routing Engines

Table 18 on page 27 lists the Routing Engine supported on the PTX1000.

NOTE: **The PTX1000 supports 64-bit Junos OS only.**

Copyright © 2018, Juniper Networks, Inc.

Table 18: PTX1000 Routing Engines

Model Number	Name in CLI Output	First Supported Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
Built-in Routing Engine	**RE-PTX1000**	• 16.1X65-D30 • 17.2R1	em0	bme0 em1

PTX3000 Routing Engines

Table 19 on page 27 lists the Routing Engines supported on the PTX3000.

NOTE: The PTX3000 supports 64-bit Junos OS only.

Table 19: PTX3000 Routing Engines

Model Number	Name in CLI Output	First Supported Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
RE-DUO-C2600-16G	**RE-DUO-2600**	13.2R2	em0	ixgbe0 ixgbe1
RCB-PTX-X6-32G	**RE-PTX-2X00x6**	16.1R4 17.1R1 This Routing Engine does not support Junos OS Release 16.2.	em0	ixlv0 ixlv1

PTX5000 Routing Engines

Table 20 on page 28 lists the Routing Engines supported on the PTX5000.

NOTE:

- PTX5000 supports 64-bit Junos OS only.
- The PTX5000 router supports two midplanes. The midplane identified as **Midplane-8S** in the CLI output is supported in Junos OS releases, 12.1X48, 12.3, and 13.2. The enhanced midplane, identified as **Midplane-8SeP** is supported from Junos OS release 14.1 onwards.

  The RE-DUO-2600 routing engine with Junos OS 13.2 or earlier is not supported on the PTX5000BASE2 midplane.

Table 20: PTX5000 Routing Engines

Model Number	Name in CLI Output	First Supported Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
RE-DUO-C2600-16G	**RE-DUO-2600**	12.1X48 12.3 13.2 NOTE: The PTX5000 does not support Junos OS Releases 12.1, 12.2, or 13.1.	em0	ixgbe0 ixgbe1
RE-PTX-X8-64G	**RE-PTX-2X00x8**	15.1F4 16.1R1	em0	ixlv0 ixlv1 em1
RE-PTX-X8-128G	**RE-PTX-2X00x8-128G**	18.1R1	em0	ixlv0 ixlv1 em1

PTX10008 and PTX10016 Routing Engines

Table 21 on page 28 lists the Routing Engines supported on the PTX10008 and PTX10016 routers.

Table 21: PTX10008 and PTX10016 Routing Engines

Model Number	Name in CLI Output	First Supported Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
JNP10K-RE0	**RE-PTX-2X00x4**	17.2R1	em0, em1	bme0 bme1
JNP10K-RE1	**JNP10K-RE1 RE X10**	18.2R!	em0	bme0 bme1

T320 Routing Engines

Table 22 on page 29 lists the Routing Engines supported by the T320 router.

 Copyright © 2018, Juniper Networks, Inc.

Table 22: T320 Routing Engines

Model Number	Name in CLI Output	First Supported 32-bit Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
RE-600-2048 (EOL details: TSB14373)	**RE-3.0** or **RE-3.0 (RE-600)**	5.3	fxp0	fxp1 fxp2
RE-1600-2048 (EOL details: TSB14374	**RE-4.0**	6.2	fxp0	fxp1 fxp2
RE-A-2000-4096	**RE-A-2000**	8.1	fxp0	fxp1 fxp2

The T320 router supports the CB-T control board.

T640 Routing Engines

Table 23 on page 29 lists the Routing Engines supported by the T640 router.

Table 23: T640 Routing Engines

Model Number	Name in CLI Output	First Supported 32-bit Junos OS Release	First Supported 64-bit Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
RE-600-2048 (EOL details: TSB14373)	**RE-3.0** or **RE-3.0 (RE-600)**	5.3	–	fxp0	fxp1 fxp2
RE-1600-2048 (EOL details: TSB14374	**RE-4.0**	6.2	–	fxp0	fxp1 fxp2
RE-A-2000-4096	**RE-A-2000**	8.1	–	fxp0	em0 bcm0
RE-DUO-C1800-8G	**RE-DUO-1800**	32-bit Junos OS on a standalone T640 router: 11.2 32-bit Junos OS on a T640 router in a routing matrix: 11.4R9	64-bit Junos OS on a standalone T640 router: 11.3 64-bit Junos OS on a T640 router in a routing matrix: 11.4R9	em0	bcm0 em1
RE-DUO-C1800-16G	**RE-DUO-1800**	32-bit Junos OS on a standalone T640 router: 11.4R2 32-bit Junos OS on a T640 router in a routing matrix: 11.4R9	64-bit Junos OS on a standalone T640 router: 11.4R2 64-bit Junos OS on a T640 router in a routing matrix: 11.4R9	em0	bcm0 em1

The T640 standalone router supports CB-T control board and CB-LCC in a T640 routing matrix.

T1600 Routing Engines

Table 24 on page 30 lists the Routing Engines supported by the T1600 router.

NOTE: (Two RE-DUO-C1800-8G or two RE-DUO-C1800-16G are required to connect to a Routing Matrix)

Table 24: T1600 Routing Engines

Model Number	Name in CLI Output	First Supported 32-bit Junos OS Release	First Supported 64-bit Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
RE-600-2048 (EOL details: TSB14373)	**RE-3.0** or **RE-3.0 (RE-600)**	8.5	–	fxp0	fxp1 fxp2
RE-1600-2048 (EOL details: TSB14374	**RE-4.0 (RE-1600)**	8.5	–	fxp0	fxp1 fxp2
RE-A-2000-4096	**RE-A-2000**	8.5	–	fxp0	em0 bcm0
RE-DUO-C1800-8G	**RE-TXP-LCC** or **RE-DUO-1800**	32-bit Junos OS on a T1600 router in a routing matrix: 9.6 NOTE: Junos OS Releases 9.6 through 10.4 support RE-DUO-C1800-8G only during upgrade to a line-card chassis (LCC) in a routing matrix. 32-bit Junos OS on a standalone T1600 router: 11.1	64-bit Junos OS on a T1600 router in a routing matrix: 9.6 64-bit Junos OS on a standalone T1600 router: 11.1	em0	bcm0 em1
RE-DUO-C1800-16G	**RE-DUO-1800**	32-bit Junos OS on a standalone T1600 router: 11.4R2 32-bit Junos OS on a T1600 router in a routing matrix: 11.4R2	64-bit Junos OS on a standalone T1600 router: 11.4R2 64-bit Junos OS on a T1600 router in a routing matrix: 11.4R2	em0	bcm0 em1

T4000 Routing Engines

Table 25 on page 31 lists the Routing Engines supported by the T4000 router.

Copyright © 2018, Juniper Networks, Inc.

NOTE: The T4000 router supports 64-bit Junos OS only.

Table 25: T4000 Routing Engines

Model Number	Name in CLI Output	First Supported 64-bit Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
RE-DUO-C1800-8G	**RE-DUO-1800**	Standalone T4000 router: 12.1 T4000 router in a routing matrix: 13.1	em0	bcm0 em1
RE-DUO-C1800-16G	**RE-DUO-1800**	Standalone T4000 router: 12.1R2 T4000 router in a routing matrix: 13.1	em0	bcm0 em1

The T4000 router supports the CB-LCC control board.

TX Matrix Routing Engines

Table 26 on page 31 lists the Routing Engines supported by the TX Matrix router.

Table 26: TX Matrix Routing Engines

Model Number	Name in CLI Output	First Supported 32-bit Junos OS Release	First Supported 64-bit Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
RE-600-2048 (EOL details: TSB14373)	**RE-3.0** or **RE-3.0 (RE-600)**	7.0	–	fxp0	fxp1 fxp2
RE-1600-2048 (EOL details: TSB14374	**RE-4.0 (RE-1600)**	7.0	–	fxp0	fxp1 fxp2
RE-A-2000-4096	**RE-A-2000**	8.5	–	fxp0	em0 bcm0
RE-DUO-C1800-8G	**RE-DUO-1800**	11.4R9	11.4R9	em0	bcm0 em1
RE-DUO-C1800-16G	**RE-DUO-1800**	11.4R9	11.4R9	em0	bcm0 em1

The TXP router supports two control boards, CB-TX and CB-LCC. The CB-LCC is required for both RE-DUO-C1800-8G and RE-DUO-C1800-16G Routing Engines.

Copyright © 2018, Juniper Networks, Inc.

TX Matrix Plus Routing Engines

Table 27 on page 32 lists the Routing Engines supported by the TX Matrix Plus router.

Table 27: TX Matrix Plus Routing Engines

Model Number	Name in CLI Output	First Supported 32-bit Junos OS Release	First Supported 64-bit Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
RE-DUO-C2600-16G	**RE-TXP-SFC** or **RE-DUO-2600**	32-bit Junos OS: 9.6	64-bit Junos OS: 11.4	em0	ixgbe0 ixgbe1

The TX Matrix Plus router supports the CB-TXP control board.

TX Matrix Plus (with 3D SIBs) Routing Engines

Table 28 on page 32 lists the Routing Engines supported by the TX Matrix Plus router with 3D SIBs.

Table 28: Routing Engines on TX Matrix Plus with 3D SIBs

Model Number	Name in CLI Output	First Supported 32-bit Junos OS Release	First Supported 64-bit Junos OS Release	Management Ethernet Interface	Internal Ethernet Interface
RE-DUO-C2600-16G	**RE-TXP-SFC** or **RE-DUO-2600**	-	64-bit Junos OS: 11.4	em0	ixgbe0 ixgbe1

Related Documentation

- *Routing Engine Specifications*
- Understanding Internal Ethernet Interfaces on page 15
- Understanding Management Ethernet Interfaces on page 10

Interface Naming Overview

Each interface has an interface name, which specifies the media type, the slot in which the FPC or DPC is located, the location on the FPC where the PIC is installed, and the PIC or DPC port. The interface name uniquely identifies an individual network connector in the system. You use the interface name when configuring interfaces and when enabling various functions and properties, such as routing protocols, on individual interfaces. The system uses the interface name when displaying information about the interface, for example, in the **show interfaces** command.

The interface name is represented by a physical part, a channel part, and a logical part in the following format:

physical<:channel>.logical

 Copyright © 2018, Juniper Networks, Inc.

The channel part of the name is optional for all interfaces except channelized DS3, E1, OC12, and STM1 interfaces.

The following sections provide interface naming configuration guidelines:

- Physical Part of an Interface Name on page 33
- Logical Part of an Interface Name on page 39
- Separators in an Interface Name on page 39
- Channel Part of an Interface Name on page 39
- Interface Naming for a Routing Matrix Based on a TX Matrix Router on page 40
- Interface Naming for a Routing Matrix Based on a TX Matrix Plus Router on page 42
- Chassis Interface Naming on page 44
- Examples: Interface Naming on page 45

Physical Part of an Interface Name

The physical part of an interface name identifies the physical device, which corresponds to a single physical network connector.

NOTE:

The internal interface is dependent on the Routing Engine. To identify if the Routing Engine is using this type of interface, use the following command:

```
user@host> show interfaces terse
Interface               Admin Link Proto    Local                 Remote
pfe-1/0/0                up    up
pfe-1/0/0.16383          up    up    inet
                                     inet6
pfh-1/0/0                up    up
pfh-1/0/0.16383          up    up    inet
[..........]
bcm0                     up    up <----------------
bcm0.0                   up    up    inet 10.0.0.1/8
[..........]
lsi                      up    up
mtun                     up    up
pimd                     up    up
pime                     up    up
tap                      up    up
```

For more information on the Routing Engines that each chassis supports, the first supported release for the Routing Engine in the specified chassis, the management Ethernet interface, and the internal Ethernet interfaces for each Routing Engine, please refer the link titled *Supported Routing Engines by Chassis* under Related Documentation section.

This part of the interface name has the following format:

type-fpc/pic/port

Copyright © 2018, Juniper Networks, Inc.

type is the media type, which identifies the network device that can be one of the following:

- **ae**—Aggregated Ethernet interface. This is a virtual aggregated link and has a different naming format from most PICs; for more information, see *Aggregated Ethernet Interfaces Overview*.
- **as**—Aggregated SONET/SDH interface. This is a virtual aggregated link and has a different naming format from most PICs; for more information, see *Configuring Aggregated SONET/SDH Interfaces*.
- **at**—ATM1 or ATM2 intelligent queuing (IQ) interface or a virtual ATM interface on a circuit emulation (CE) interface.
- **bcm**—The bcm0 internal Ethernet process is supported on specific Routing engines for various M series and T series routers. For more information please refer the link titled *Supported Routing Engines by Chassis* under Related Documentation section.
- **cau4**—Channelized AU-4 IQ interface (configured on the Channelized STM1 IQ or IQE PIC or Channelized OC12 IQ and IQE PICs).
- **ce1**—Channelized E1 IQ interface (configured on the Channelized E1 IQ PIC or Channelized STM1 IQ or IQE PIC).
- **ci**—Container interface.
- **coc1**—Channelized OC1 IQ interface (configured on the Channelized OC12 IQ and IQE or Channelized OC3 IQ and IQE PICs).
- **coc3**—Channelized OC3 IQ interface (configured on the Channelized OC3 IQ and IQE PICs).
- **coc12**—Channelized OC12 IQ interface (configured on the Channelized OC12 IQ and IQE PICs).
- **coc48**—Channelized OC48 interface (configured on the Channelized OC48 and Channelized OC48 IQE PICs).
- **cp**—Collector interface (configured on the Monitoring Services II PIC).
- **cstm1**—Channelized STM1 IQ interface (configured on the Channelized STM1 IQ or IQE PIC).
- **cstm4**—Channelized STM4 IQ interface (configured on the Channelized OC12 IQ and IQE PICs).
- **cstm16**—Channelized STM16 IQ interface (configured on the Channelized OC48/STM16 and Channelized OC48/STM16 IQE PICs).
- **ct1**—Channelized T1 IQ interface (configured on the Channelized DS3 IQ and IQE PICs, Channelized OC3 IQ and IQE PICs, Channelized OC12 IQ and IQE PICs, or Channelized T1 IQ PIC).
- **ct3**—Channelized T3 IQ interface (configured on the Channelized DS3 IQ and IQE PICs, Channelized OC3 IQ and IQE PICs, or Channelized OC12 IQ and IQE PICs).
- **demux**—Interface that supports logical IP interfaces that use the IP source or destination address to demultiplex received packets. Only one demux interface (**demux0**) exists

 Copyright © 2018, Juniper Networks, Inc.

per chassis. All demux logical interfaces must be associated with an underlying logical interface.

- **dfc**—Interface that supports dynamic flow capture processing on T Series or M320 routers containing one or more Monitoring Services III PICs. Dynamic flow capture enables you to capture packet flows on the basis of dynamic filtering criteria. Specifically, you can use this feature to forward passively monitored packet flows that match a particular filter list to one or more destinations using an on-demand control protocol.
- **ds**—DS0 interface (configured on the Multichannel DS3 PIC, Channelized E1 PIC, Channelized OC3 IQ and IQE PICs, Channelized OC12 IQ and IQE PICs, Channelized DS3 IQ and IQE PICs, Channelized E1 IQ PIC, Channelized STM1 IQ or IQE PIC, or Channelized T1 IQ).
- **dsc**—Discard interface.
- **e1**—E1 interface (including channelized STM1-to-E1 interfaces).
- **e3**—E3 interface (including E3 IQ interfaces).
- **em**—Management and internal Ethernet interfaces. For M Series routers, MX Series routers, T Series routers, and TX Series routers, you can use the **show chassis hardware** command to display hardware information about the router, including its Routing Engine model. To determine which management interface is supported on your router and Routing Engine combination, see "Understanding Management Ethernet Interfaces" on page 10 and "Supported Routing Engines by Router" on page 17.
- **es**—Encryption interface.
- **et**—100-Gigabit Ethernet interfaces (10, 40, and 100-Gigabit Ethernet interface for PTX Series Packet Transport Routers only).
- **fe**—Fast Ethernet interface.
- **fxp**—Management and internal Ethernet interfaces. For M Series routers, MX Series routers, T Series routers, and TX Series routers, you can use the **show chassis hardware** command to display hardware information about the router, including its Routing Engine model. To determine which management interface is supported on your router and Routing Engine combination, see "Understanding Management Ethernet Interfaces" on page 10 and "Supported Routing Engines by Router" on page 17.
- **ge**—Gigabit Ethernet interface.

NOTE:

- The XENPAK 10-Gigabit Ethernet interface PIC, which is supported only on M series routers, is configured using the **ge** interface naming convention instead of the **xe** interface naming convention. Refer the following show commands for more information:

```
user@host> show chassis hardware
..
FPC 4            REV 02   710-015839   CZ1853           M120 FPC
Type 3
  PIC 0          REV 09   750-009567   NH1857           1x
10GE(LAN),XENPAK
    Xcvr 0       REV 01   740-012045   535TFZX6         XENPAK-SR

user@host> show configuration interfaces ge-4/0/0
unit 0 {
    family inet {
        address 100.0.0.1/24;
    }
}
```

- In MX and SRX series devices, the 1 and 10-Gigabit SFP or SFP+ optical interfaces are always named as **xe** even if a 1-Gigabit SFP is inserted. However, in EX and QFX series devices, the interface name is shown as **ge** or **xe** based on the speed of the optical device inserted.

- **gr**—Generic routing encapsulation (GRE) tunnel interface.
- **gre**—Internally generated interface that is configurable only as the control channel for Generalized MPLS (GMPLS). For more information about GMPLS, see the *Junos OS MPLS Applications Library for Routing Devices*.

NOTE: You can configure GRE interfaces (gre-x/y/z) only for GMPLS control channels. GRE interfaces are not supported or configurable for other applications..

- **ip**—IP-over-IP encapsulation tunnel interface.
- **ipip**—Internally generated interface that is not configurable.
- **ixgbe**—The internal Ethernet process ixgbe0 and ixgbe1 are used by the RE-DUO-C2600-16G Routing Engine, which is supported on TX Matrix Plus and PTX5000.
- **iw**—Logical interfaces associated with the endpoints of Layer 2 circuit and Layer 2 VPN connections (pseudowire stitching Layer 2 VPNs). For more information about VPNs, see the *Junos OS VPNs Library for Routing Devices*.
- **lc**—Internally generated interface that is not configurable.

Copyright © 2018, Juniper Networks, Inc.

- **lo**—Loopback interface. The Junos OS automatically configures one loopback interface (**lo0**). The logical interface **lo0.16383** is a nonconfigurable interface for router control traffic.
- **ls**—Link services interface.
- **lsi**—Internally generated interface that is not configurable.
- **ml**—Multilink interface (including Multilink Frame Relay and MLPPP).
- **mo**—Monitoring services interface (including monitoring services and monitoring services II). The logical interface **mo-*fpc/pic/port*.16383** is an internally generated, nonconfigurable interface for router control traffic.
- **ms**—Multiservices interface.
- **mt**—Multicast tunnel interface (internal router interface for VPNs). If your router has a Tunnel PIC, the Junos OS automatically configures one multicast tunnel interface (**mt**) for each virtual private network (VPN) you configure. Although it is not necessary to configure multicast interfaces, you can use the **multicast-only** statement to configure the unit and family so that the tunnel can transmit and receive multicast traffic only. For more information, see multicast-only.
- **mtun**—Internally generated interface that is not configurable.
- **oc3**—OC3 IQ interface (configured on the Channelized OC12 IQ and IQE PICs or Channelized OC3 IQ and IQE PICs).
- **pd**—Interface on the rendezvous point (RP) that de-encapsulates packets.
- **pe**—Interface on the first-hop PIM router that encapsulates packets destined for the RP router.
- **pimd**—Internally generated interface that is not configurable.
- **pime**—Internally generated interface that is not configurable.
- **rlsq**—Container interface, numbered from 0 through 127, used to tie the primary and secondary LSQ PICs together in high availability configurations. Any failure of the primary PIC results in a switch to the secondary PIC and vice versa.
- **rms**—Redundant interface for two multiservices interfaces.
- **rsp**—Redundant virtual interface for the adaptive services interface.
- **se**—Serial interface (including EIA-530, V.35, and X.21 interfaces).
- **si**—Services-inline interface, which is hosted on a Trio-based line card.
- **so**—SONET/SDH interface.
- **sp**—Adaptive services interface. The logical interface **sp-*fpc/pic/port*.16383** is an internally generated, nonconfigurable interface for router control traffic.
- **stm1**—STM1 interface (configured on the OC3/STM1 interfaces).
- **stm4**—STM4 interface (configured on the OC12/STM4 interfaces).
- **stm16**—STM16 interface (configured on the OC48/STM16 interfaces).

Copyright © 2018, Juniper Networks, Inc.

- **t1**—T1 interface (including channelized DS3-to-DS1 interfaces).
- **t3**—T3 interface (including channelized OC12-to-DS3 interfaces).
- **tap**—Internally generated interface that is not configurable.
- **umd**—USB modem interface.
- **vsp**—Voice services interface.
- **vc4**—Virtually concatenated interface.
- **vt**—Virtual loopback tunnel interface.
- **xe**—10-Gigabit Ethernet interface. Some older 10-Gigabit Ethernet interfaces use the **ge** media type (rather than **xe**) to identify the physical part of the network device.
- **xt**—Logical interface for Protected System Domains to establish a Layer 2 tunnel connection.

fpc identifies the number of the FPC or DPC card on which the physical interface is located. Specifically, it is the number of the slot in which the card is installed.

M40, M40e, M160, M320, M120, T320, T640, and T1600 routers each have eight FPC slots that are numbered 0 through 7, from left to right as you are facing the front of the chassis. For information about compatible FPCs and PICs, see the hardware guide for your router.

On PTX1000 routers, the FPC number is always 0.

The M20 router has four FPC slots that are numbered 0 through 3, from top to bottom as you are facing the front of the chassis. The slot number is printed adjacent to each slot.

MX Series routers support DPCs, FPCs, and Modular Interface Cards (MICs). For information about compatible DPCs, FPCs, PICs, and MICs, see the *MX Series Interface Module Reference*.

For M5, M7i, M10, and M10i routers, the FPCs are built into the chassis; you install the PICs into the chassis.

The M5 and M7i routers have space for up to four PICs. The M7i router also comes with an integrated Tunnel PIC, or an optional integrated AS PIC, or an optional integrated MS PIC.

The M10 and M10i routers have space for up to eight PICs.

A routing matrix can have up to 32 FPCs (numbered 0 through 31).

For more information about interface naming for a routing matrix, see "Interface Naming for a Routing Matrix Based on a TX Matrix Router" on page 40.

pic identifies the number of the PIC on which the physical interface is located. Specifically, it is the number of the PIC location on the FPC. FPCs with four PIC slots are numbered 0 through 3. FPCs with three PIC slots are numbered 0 through 2. The PIC location is printed

 Copyright © 2018, Juniper Networks, Inc.

on the FPC carrier board. For PICs that occupy more than one PIC slot, the lower PIC slot number identifies the PIC location.

port identifies a specific port on a PIC or DPC. The number of ports varies depending on the PIC. The port numbers are printed on the PIC.

Logical Part of an Interface Name

The logical unit part of the interface name corresponds to the logical unit number, which can be a number from 0 through 16,385 for all interface types except demux and PPPoE. For these two interface types only, the range is 0 through 65,535.

In the virtual part of the name, a period (**.**) separates the port and logical unit numbers:

- Other platforms:

  type-fpc/pic/**port.*****logical***

Separators in an Interface Name

There is a separator between each element of an interface name.

In the physical part of the name, a hyphen (-) separates the media type from the FPC number, and a slash (/) separates the FPC, PIC, and port numbers.

In the virtual part of the name, a period (.) separates the channel and logical unit numbers.

A colon (:) separates the physical and virtual parts of the interface name.

Channel Part of an Interface Name

The channel identifier part of the interface name is required only on channelized interfaces. For channelized interfaces, channel 0 identifies the first channelized interface. For channelized IQ and channelized IQE interfaces, channel 1 identifies the first channelized interface. A nonconcatenated (that is, channelized) SONET/SDH OC48 interface has four OC12 channels, numbered 0 through 3.

To determine which types of channelized PICs are currently installed in the router, use the **show chassis hardware** command from the top level of the command-line interface (CLI). Channelized IQ and IQE PICs are listed in the output with “intelligent queuing IQ” or “enhanced intelligent queuing IQE” in the description. For more information, see *Channelized Interfaces Overview*.

For ISDN interfaces, you specify the B-channel in the form **bc-*pim*/0/*port*:*n***. ***n*** is the B-channel ID and can be 1 or 2. You specify the D-channel in the form **dc-*pim*/0/*port*:0**.

NOTE: **For ISDN, the B-channel and D-channel interfaces do not have any configurable parameters. However, when interface statistics are displayed, B-channel and D-channel interfaces have statistical values.**

NOTE: In the Junos OS implementation, the term *logical interfaces* generally refers to interfaces you configure by including the **unit** statement at the **[edit interfaces *interface-name*]** hierarchy level. Logical interfaces have the *.logical* descriptor at the end of the interface name, as in **ge-0/0/0.1** or **t1-0/0/0:0.1**, where the logical unit number is 1.

Although channelized interfaces are generally thought of as logical or virtual, the Junos OS sees T3, T1, and NxDS0 interfaces within a channelized IQ or IQE PIC as physical interfaces. For example, both **t3-0/0/0** and **t3-0/0/0:1** are treated as physical interfaces by the Junos OS. In contrast, **t3-0/0/0.2** and **t3-0/0/0:1.2** are considered logical interfaces because they have the .2 at the end of the interface names.

Interface Naming for a Routing Matrix Based on a TX Matrix Router

A routing matrix based on a Juniper Networks TX Matrix router is a multichassis architecture composed of one TX Matrix router and from one to four interconnected T640 routers. From the perspective of the user interface, the routing matrix appears as a single router. The TX Matrix router controls all the T640 routers, as shown in Figure 3 on page 40.

Figure 3: Routing Matrix

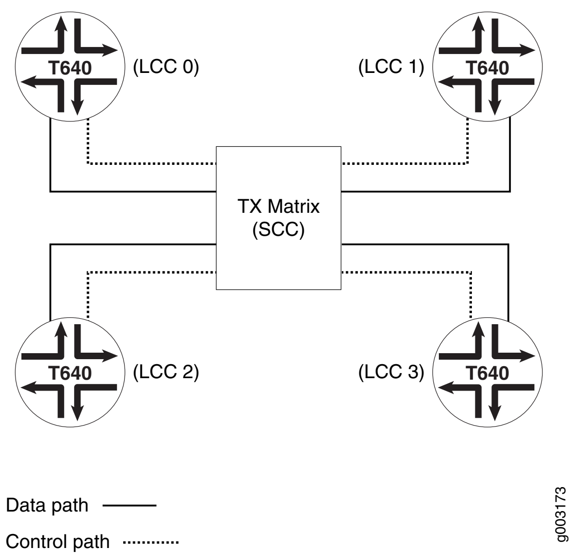

A TX Matrix router is also referred to as a *switch-card chassis* (SCC). The CLI uses **scc** to refer to the TX Matrix router. A T640 router in a routing matrix is also referred to as a *line-card chassis* (LCC). The CLI uses **lcc** as a prefix to refer to a specific T640 router.

LCCs are assigned numbers 0 through 3, depending on the hardware setup and connectivity to the TX Matrix router. For more information, see the *TX Matrix Router Hardware Guide*. A routing matrix can have up to four T640 routers, and each T640 router has up to eight FPCs. Therefore, the routing matrix as a whole can have up to 32 FPCs (0 through 31).

 Copyright © 2018, Juniper Networks, Inc.

In the Junos OS CLI, an interface name has the following format:

type-fpc/pic/port

When you specify the ***fpc*** number for a T640 router in a routing matrix, the Junos OS determines which T640 router contains the specified FPC based on the following assignment:

- On LCC 0, FPC hardware slots 0 through 7 are configured as 0 through 7.
- On LCC 1, FPC hardware slots 0 through 7 are configured as 8 through 15.
- On LCC 2, FPC hardware slots 0 through 7 are configured as 16 through 23.
- On LCC 3, FPC hardware slots 0 through 7 are configured as 24 through 31.

For example, the **1** in **se-1/0/0** refers to FPC hardware slot 1 on the T640 router labeled **lcc0**. The **11** in **t1-11/2/0** refers to FPC hardware slot 3 on the T640 router labeled **lcc1**. The **20** in **so-20/0/1** refers to FPC hardware slot 4 on the T640 router labeled **lcc2**. The **31** in **t3-31/1/0** refers to FPC hardware slot 7 on the T640 router labeled **lcc3**.

Table 29 on page 41 summarizes the FPC numbering for a T640 router in a routing matrix.

Table 29: FPC Numbering for T640 Routers in a Routing Matrix

LCC Numbers Assigned to the T640 Router	Configuration Numbers
0	0 through 7
1	8 through 15
2	16 through 23
3	24 through 31

Table 30 on page 41 lists each FPC hardware slot and the corresponding configuration numbers for LCCs 0 through 3.

Table 30: One-to-One FPC Numbering for T640 Routers in a Routing Matrix

FPC Numbering	T640 Routers							
	LCC 0							
Hardware Slots	0	1	2	3	4	5	6	7
Configuration Numbers	0	1	2	3	4	5	6	7
	LCC 1							
Hardware Slots	0	1	2	3	4	5	6	7

Copyright © 2018, Juniper Networks, Inc.

Table 30: One-to-One FPC Numbering for T640 Routers in a Routing Matrix (continued)

FPC Numbering	T640 Routers							
Configuration Numbers	8	9	10	11	12	13	14	15
	LCC 2							
Hardware Slots	0	1	2	3	4	5	6	7
Configuration Numbers	16	17	18	19	20	21	22	23
	LCC 3							
Hardware Slots	0	1	2	3	4	5	6	7
Configuration Numbers	24	25	26	27	28	29	30	31

Interface Naming for a Routing Matrix Based on a TX Matrix Plus Router

A routing matrix based on a Juniper Networks TX Matrix Plus Router is a multichassis architecture composed of one TX Matrix Plus router and from one to four interconnected T1600 routers. From the perspective of the user interface, the routing matrix appears as a single router. The TX Matrix Plus router controls all the T1600 routers, as shown in Figure 4 on page 42.

Figure 4: Routing Matrix Based on a TX Matrix Plus Router

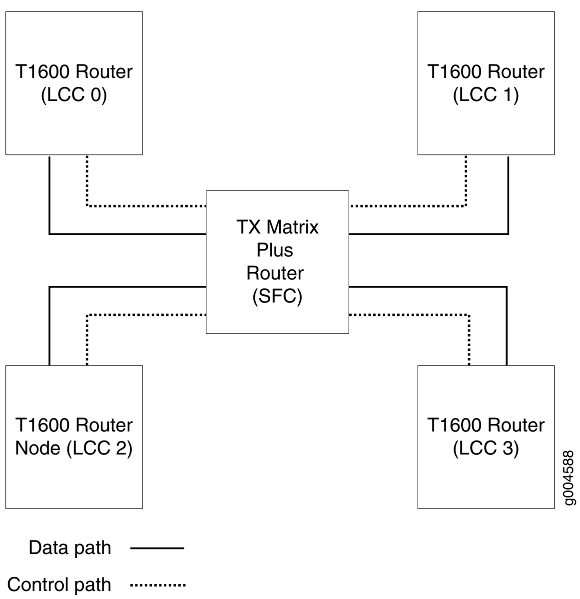

A TX Matrix Plus router is also referred to as a *switch-fabric chassis* (SFC). The CLI uses **sfc** to refer to the TX Matrix Plus router. A T1600 router in a routing matrix is also referred to as a *line-card chassis* (LCC). The CLI uses **lcc** as a prefix to refer to a specific T1600 router.

 Copyright © 2018, Juniper Networks, Inc.

LCCs are assigned numbers, 0 through 3, depending on the hardware setup and connectivity to the TX Matrix Plus router. For more information, see the *TX Matrix Plus Router Hardware Guide*. A routing matrix based on a TX Matrix Plus router can have up to four T1600 routers, and each T1600 router has up to eight FPCs. Therefore, the routing matrix as a whole can have up to 32 FPCs (0 through 31).

In the Junos OS CLI, an interface name has the following format:

type-fpc/pic/port

When you specify the ***fpc*** number for a T1600 router in a routing matrix, the Junos OS determines which T1600 router contains the specified FPC based on the following assignment:

- On LCC 0, FPC hardware slots 0 through 7 are configured as 0 through 7.
- On LCC 1, FPC hardware slots 0 through 7 are configured as 8 through 15.
- On LCC 2, FPC hardware slots 0 through 7 are configured as 16 through 23.
- On LCC 3, FPC hardware slots 0 through 7 are configured as 24 through 31.

For example, the **1** in **se-1/0/0** refers to FPC hardware slot 1 on the T1600 router labeled **lcc0**. The **11** in **t1-11/2/0** refers to FPC hardware slot 3 on the T1600 router labeled **lcc1**. The **20** in **so-20/0/1** refers to FPC hardware slot 4 on the T1600 router labeled **lcc2**. The **31** in **t3-31/1/0** refers to FPC hardware slot 7 on the T1600 router labeled **lcc3**.

Table 31 on page 43 summarizes the FPC numbering for a routing matrix based on a TX Matrix Plus router.

Table 31: FPC Numbering for T1600 Routers in a Routing Matrix

LCC Numbers Assigned to the T1600 Router	Configuration Numbers
0	0 through 7
1	8 through 15
2	16 through 23
3	24 through 31

Table 32 on page 43 lists each FPC hardware slot and the corresponding configuration numbers for LCCs 0 through 3.

Table 32: One-to-One FPC Numbering for T1600 Routers in a Routing Matrix

FPC Numbering	T1600 Routers							
	LCC 0							
Hardware Slots	0	1	2	3	4	5	6	7

Copyright © 2018, Juniper Networks, Inc.

Table 32: One-to-One FPC Numbering for T1600 Routers in a Routing Matrix (continued)

FPC Numbering	T1600 Routers							
Configuration Numbers	0	1	2	3	4	5	6	7
	LCC 1							
Hardware Slots	0	1	2	3	4	5	6	7
Configuration Numbers	8	9	10	11	12	13	14	15
	LCC 2							
Hardware Slots	0	1	2	3	4	5	6	7
Configuration Numbers	16	17	18	19	20	21	22	23
	LCC 3							
Hardware Slots	0	1	2	3	4	5	6	7
Configuration Numbers	24	25	26	27	28	29	30	31

Chassis Interface Naming

You configure some PIC properties, such as framing, at the **[edit chassis]** hierarchy level. Chassis interface naming varies depending on the routing hardware.

- To configure PIC properties for a standalone router, you must specify the FPC and PIC numbers, as follows:

```
[edit chassis]
fpc slot-number {
  pic pic-number {
    ...
  }
}
```

- To configure PIC properties for a T640 or T1600 router configured in a routing matrix, you must specify the LCC, FPC, and PIC numbers, as follows:

```
[edit chassis]
lcc lcc-number {
  fpc slot-number { # Use the hardware FPC slot number
    pic pic-number {
      ...
    }
  }
}
```

  For the FPC slot in a T640 router in a routing matrix, specify the actual hardware slot number, as labeled on the T640 router chassis. Do not use the corresponding software FPC configuration numbers shown in Table 30 on page 41.

 Copyright © 2018, Juniper Networks, Inc.

For the FPC slot in a T1600 router in a routing matrix, specify the actual hardware slot number, as labeled on the T1600 router chassis. Do not use the corresponding software FPC configuration numbers shown in Table 31 on page 43.

For more information about the **[edit chassis]** hierarchy, see the *Junos OS Administration Library*.

Examples: Interface Naming

This section provides examples of naming interfaces. For an illustration of where slots, PICs, and ports are located, see Figure 5 on page 45.

Figure 5: Interface Slot, PIC, and Port Locations

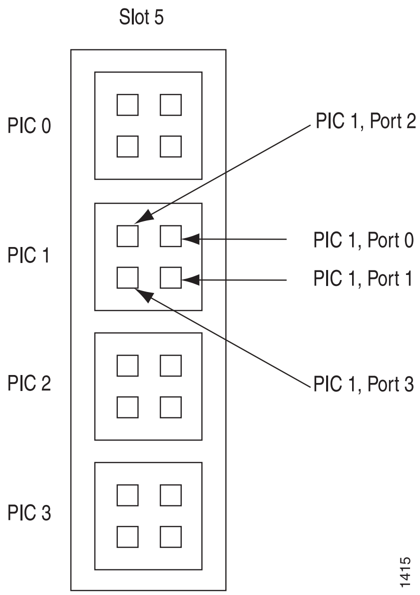

For an FPC in slot 1 with two OC3 SONET/SDH PICs in PIC positions 0 and 1, each PIC with two ports uses the following names:

```
so-1/0/0.0
so-1/0/1.0
so-1/1/0.0
so-1/1/1.0
```

An OC48 SONET/SDH PIC in slot 1 and in concatenated mode appears as a single FPC with a single PIC, which has a single port. If this interface has a single logical unit, it has the following name:

```
so-1/0/0.0
```

An OC48 SONET/SDH PIC in slot 1 and in channelized mode has a number for each channel. For example:

```
so-1/0/0:0
so-1/0/0:1
```

For an FPC in slot 1 with a Channelized OC12 PIC in PIC position 2, the DS3 channels have the following names:

Copyright © 2018, Juniper Networks, Inc.

t3-1/2/0:0
t3-1/2/0:1
t3-1/2/0:2
...
t3-1/2/0:11

For an FPC in slot 1 with four OC12 ATM PICs (the FPC is fully populated), the four PICs, each with a single port and a single logical unit, have the following names:

at-1/0/0.0
at-1/1/0.0
at-1/2/0.0
at-1/3/0.0

In a routing matrix on the T640 router labeled **lcc1**, for an FPC in slot 5 with four SONET OC192 PICs, the four PICs, each with a single port and a single logical unit, have the following names:

so-13/0/0.0
so-13/1/0.0
so-13/2/0.0
so-13/3/0.0

For an FPC in slot 1 with one 4-port ISDN BRI interface card, port 4 has the following name:

br-1/0/4

The first B-channel, the second B-channel, and the control channel have the following names:

bc-1/0/4:1
bc-1/0/4:2
dc-1/0/4:0

Related Documentation

- Router Interfaces Overview on page 4
- Physical Part of an Interface Name on page 59

- Supported Routing Engines by Chassis

Interface Encapsulations Overview

Table 33 on page 47 lists encapsulation support by interface type.

Copyright © 2018, Juniper Networks, Inc.

Table 33: Encapsulation Support by Interface Type

Interface Type	Physical Interface Encapsulation	Logical Interface Encapsulation
ae—Aggregated Ethernet interface	**ethernet-ccc**—Ethernet cross-connect **extended-vlan-ccc**—Nonstandard TPID tagging for a cross-connect **extended-vlan-vpls**—Extended VLAN virtual private LAN service **flexible-ethernet-services**—Allows per-unit Ethernet encapsulation configuration **vlan-ccc**—802.1Q tagging for a cross-connect **ethernet-vpls**—Ethernet virtual private LAN service **vlan-vpls**—VLAN virtual private LAN service	**dix**—Ethernet DIXv2 (RFC 894) **vlan-ccc**—802.1Q tagging for a cross-connect
as—Aggregated SONET/SDH interface	**cisco-hdlc**—Cisco-compatible HDLC framing **ppp**—Serial PPP device	NA
at—ATM1 interface	**atm-ccc-cell-relay**—ATM cell relay encapsulation for a cross-connect **atm-pvc**—ATM permanent virtual circuits **ethernet-over-atm**—Ethernet over ATM encapsulation	**atm-ccc-cell-relay**—ATM cell relay for CCC **atm-ccc-vc-mux**—ATM VC for CCC **atm-cisco-nlpid**—Cisco-compatible ATM NLPID encapsulation **atm-nlpid**—ATM NLPID encapsulation **atm-snap**—ATM LLC/SNAP encapsulation **atm-tcc-snap**—ATM LLC/SNAP for a translational cross-connect **atm-tcc-vc-mux**—ATM VC for a translational cross-connect **atm-vc-mux**—ATM VC multiplexing **ether-over-atm-llc**—Ethernet over ATM (LLC/SNAP) encapsulation

Copyright © 2018, Juniper Networks, Inc.

Table 33: Encapsulation Support by Interface Type (continued)

Interface Type	Physical Interface Encapsulation	Logical Interface Encapsulation
at—ATM2 intelligent queuing (IQ) interface	**atm-ccc-cell-relay**—ATM cell relay encapsulation for a cross-connect **atm-pvc**—ATM permanent virtual circuits **ethernet-over-atm**—Ethernet over ATM encapsulation	**atm-ccc-cell-relay**—ATM cell relay for CCC **atm-ccc-vc-mux**—ATM VC for CCC **atm-cisco-nlpid**—Cisco-compatible ATM NLPID encapsulation **atm-mlppp-llc**—ATM MLPPP over AAL5/LLC **atm-nlpid**—ATM NLPID encapsulation **atm-ppp-llc**—ATM PPP over AAL5/LLC **atm-ppp-vc-mux**—ATM PPP over raw AAL5 **atm-snap**—ATM LLC/SNAP encapsulation **atm-tcc-snap**—ATM LLC/SNAP for a translational cross-connect **atm-tcc-vc-mux**—ATM VC for a translational cross-connect **atm-vc-mux**—ATM VC multiplexing **ether-over-atm-llc**—Ethernet over ATM (LLC/SNAP) encapsulation **ether-vpls-over-atm-llc**—Ethernet VPLS over ATM (bridging) encapsulation
bcm—Gigabit Ethernet internal interfaces	NA	NA
br—Integrated Services Digital Network (ISDN) interface	NA	NA
ci—Container interface	**cisco-hdlc**—Cisco-compatible HDLC framing **ppp**—Serial PPP device	**aps**—SONET interface required for APS configuration.

 Copyright © 2018, Juniper Networks, Inc.

Table 33: Encapsulation Support by Interface Type (continued)

Interface Type	Physical Interface Encapsulation	Logical Interface Encapsulation
ds—DS0 interface	**cisco-hdlc**—Cisco-compatible HDLC framing **cisco-hdlc-ccc**—Cisco-compatible HDLC framing for a cross-connect **cisco-hdlc-tcc**—Cisco-compatible HDLC framing for a translational cross-connect **extended-frame-relay-ccc**—Any Frame Relay DLCI for a cross-connect **extended-frame-relay-tcc**—Any Frame Relay DLCI for a translational cross-connect **flexible-frame-relay**—Multiple Frame Relay encapsulations **frame-relay**—Frame Relay encapsulation **frame-relay-ccc**—Frame Relay for a cross-connect **frame-relay-port-ccc**—Frame Relay port encapsulation for a cross-connect **frame-relay-tcc**—Frame Relay for a translational cross-connect **multilink-frame-relay-uni-nni**—Multilink Frame Relay UNI NNI (FRF.16) encapsulation **ppp**—Serial PPP device **ppp-ccc**—Serial PPP device for a cross-connect **ppp-tcc**—Serial PPP device for a translational cross-connect	**frame-relay-ccc**—Frame Relay DLCI for CCC **frame-relay-ppp**—PPP over Frame Relay **frame-relay-tcc**—Frame Relay DLCI for a translational cross-connect
dsc—Discard interface	NA	NA

Copyright © 2018, Juniper Networks, Inc.

Table 33: Encapsulation Support by Interface Type (continued)

Interface Type	Physical Interface Encapsulation	Logical Interface Encapsulation
e1—E1 interface (including channelized STM1-to-E1 interfaces)	**cisco-hdlc**—Cisco-compatible HDLC framing **cisco-hdlc-ccc**—Cisco-compatible HDLC framing for a cross-connect **cisco-hdlc-tcc**—Cisco-compatible HDLC framing for a translational cross-connect **extended-frame-relay-ccc**—Any Frame Relay DLCI for a cross-connect **extended-frame-relay-tcc**—Any Frame Relay DLCI for a translational cross-connect **flexible-frame-relay**—Multiple Frame Relay encapsulations **frame-relay**—Frame Relay encapsulation **frame-relay-ccc**—Frame Relay for a cross-connect **frame-relay-port-ccc**—Frame Relay port encapsulation for a cross-connect **frame-relay-tcc**—Frame Relay for a translational cross-connect **multilink-frame-relay-uni-nni**—Multilink Frame Relay UNI NNI (FRF.16) encapsulation **ppp**—Serial PPP device **ppp-ccc**—Serial PPP device for a cross-connect **ppp-tcc**—Serial PPP device for a translational cross-connect	**frame-relay-ccc**—Frame Relay DLCI for CCC **frame-relay-ppp**—PPP over Frame Relay **frame-relay-tcc**—Frame Relay DLCI for a translational cross-connect

 Copyright © 2018, Juniper Networks, Inc.

Table 33: Encapsulation Support by Interface Type (continued)

Interface Type	Physical Interface Encapsulation	Logical Interface Encapsulation
e3—E3 interface (including E3 IQ and IQE interfaces)	**cisco-hdlc**—Cisco-compatible HDLC framing **cisco-hdlc-ccc**—Cisco-compatible HDLC framing for a cross-connect **cisco-hdlc-tcc**—Cisco-compatible HDLC framing for a translational cross-connect **extended-frame-relay-ccc**—Any Frame Relay DLCI for a cross-connect **extended-frame-relay-tcc**—Any Frame Relay DLCI for a translational cross-connect **flexible-frame-relay**—Multiple Frame Relay encapsulations **frame-relay**—Frame Relay encapsulation **frame-relay-ccc**—Frame Relay for a cross-connect **frame-relay-port-ccc**—Frame Relay port encapsulation for a cross-connect **frame-relay-tcc**—Frame Relay for a translational cross-connect **ppp**—Serial PPP device **ppp-ccc**—Serial PPP device for a cross-connect **ppp-tcc**—Serial PPP device for a translational cross-connect	**frame-relay-ccc**—Frame Relay DLCI for CCC **frame-relay-ppp**—PPP over Frame Relay **frame-relay-tcc**—Frame Relay DLCI for a translational cross-connect
em—Management and internal Ethernet interfaces	NA	NA

Copyright © 2018, Juniper Networks, Inc.

Table 33: Encapsulation Support by Interface Type (continued)

Interface Type	Physical Interface Encapsulation	Logical Interface Encapsulation
fe—Fast Ethernet interface	**ethernet-ccc**—Ethernet cross-connect **ethernet-tcc**—Ethernet translational cross-connect **ethernet-vpls**—Ethernet virtual private LAN service **extended-vlan-ccc**—Nonstandard TPID tagging for a cross-connect **extended-vlan-tcc**—802.1Q tagging for a translational cross-connect **extended-vlan-vpls**—Extended VLAN virtual private LAN service **vlan-ccc**—802.1Q tagging for a cross-connect **vlan-vpls**—VLAN virtual private LAN service	**dix**—Ethernet DIXv2 (RFC 894) **vlan-ccc**—802.1Q tagging for a cross-connect **vlan-vpls**—VLAN virtual private LAN service
fxp—Management and internal Ethernet interfaces	NA	NA
ge—Gigabit Ethernet interface (including Gigabit Ethernet IQ interfaces)	**ethernet-ccc**—Ethernet cross-connect **ethernet-tcc**—Ethernet translational cross-connect **ethernet-vpls**—Ethernet virtual private LAN service **extended-vlan-ccc**—Nonstandard TPID tagging for a cross-connect **extended-vlan-tcc**—802.1Q tagging for a translational cross-connect **extended-vlan-vpls**—Extended VLAN virtual private LAN service **flexible-ethernet-services**—Allows per-unit Ethernet encapsulation configuration **vlan-ccc**—802.1Q tagging for a cross-connect **vlan-vpls**—VLAN virtual private LAN service	**dix**—Ethernet DIXv2 (RFC 894) **vlan-ccc**—802.1Q tagging for a cross-connect **vlan-tcc**—802.1Q tagging for a translational cross-connect **vlan-vpls**—VLAN virtual private LAN service
ixgbe—10-Gigabit Ethernet internal interfaces	NA	NA
lo—Loopback interface; the Junos OS automatically configures one loopback interface (**lo0**)	NA	NA

Copyright © 2018, Juniper Networks, Inc.

Table 33: Encapsulation Support by Interface Type (continued)

Interface Type	Physical Interface Encapsulation	Logical Interface Encapsulation
ls—Link services interface	**multilink-frame-relay-uni-nni**—Multilink Frame Relay UNI NNI (FRF.16) encapsulation	**multilink-frame-relay-end-to-end**—Multilink Frame Relay end-to-end (FRF.15) **multilink-ppp**—Multilink PPP
lsq—Link services IQ interface	**multilink-frame-relay-uni-nni**—Multilink Frame Relay UNI NNI (FRF.16) encapsulation	**multilink-frame-relay-end-to-end**—Multilink Frame Relay end-to-end (FRF.15) **multilink-ppp**—Multilink PPP
lt—Logical tunnel interface	NA	**ethernet**—Ethernet service **ethernet-vpls**—Ethernet virtual private LAN service **ethernet-ccc**—Ethernet cross-connect **frame-relay**—Frame Relay encapsulation **frame-relay-ccc**—Frame Relay for a cross-connect **vlan**—VLAN service **vlan-ccc**—802.1Q tagging for a cross-connect **vlan-vpls**—VLAN virtual private LAN service
ml—Multilink interface (including Multilink Frame Relay and MLPPP)	NA	**multilink-frame-relay-end-to-end**—Multilink Frame Relay end-to-end (FRF.15) **multilink-ppp**—Multilink PPP

Copyright © 2018, Juniper Networks, Inc.

Table 33: Encapsulation Support by Interface Type (continued)

Interface Type	Physical Interface Encapsulation	Logical Interface Encapsulation
se—Serial interface (including EIA-530, V.35, and X.21 interfaces)	**cisco-hdlc**—Cisco-compatible HDLC framing **cisco-hdlc-ccc**—Cisco-compatible HDLC framing for a cross-connect **cisco-hdlc-tcc**—Cisco-compatible HDLC framing for a translational cross-connect **frame-relay**—Frame Relay encapsulation **frame-relay-ccc**—Frame Relay for a cross-connect **frame-relay-port-ccc**—Frame Relay port encapsulation for a cross-connect **frame-relay-tcc**—Frame Relay for a translational cross-connect **ppp**—Serial PPP device **ppp-ccc**—Serial PPP device for a cross-connect **ppp-tcc**—Serial PPP device for a translational cross-connect	**frame-relay-ccc**—Frame Relay DLCI for CCC **frame-relay-ppp**—PPP over Frame Relay **frame-relay-tcc**—Frame Relay DLCI for a translational cross-connect

Copyright © 2018, Juniper Networks, Inc.

Table 33: Encapsulation Support by Interface Type (continued)

Interface Type	Physical Interface Encapsulation	Logical Interface Encapsulation
so—SONET/SDH interface	**cisco-hdlc**—Cisco-compatible HDLC framing **cisco-hdlc-ccc**—Cisco-compatible HDLC framing for a cross-connect **cisco-hdlc-tcc**—Cisco-compatible HDLC framing for a translational cross-connect **extended-frame-relay-ccc**—Any Frame Relay DLCI for a cross-connect **extended-frame-relay-tcc**—Any Frame Relay DLCI for a translational cross-connect **flexible-frame-relay**—Multiple Frame Relay encapsulations **frame-relay**—Frame Relay encapsulation **frame-relay-ccc**—Frame Relay for a cross-connect **frame-relay-port-ccc**—Frame Relay port encapsulation for a cross-connect **frame-relay-tcc**—Frame Relay for a translational cross-connect **ppp**—Serial PPP device **ppp-ccc**—Serial PPP device for a cross-connect **ppp-tcc**—Serial PPP device for a translational cross-connect	**frame-relay-ccc**—Frame Relay DLCI for CCC **frame-relay-ppp**—PPP over Frame Relay **frame-relay-tcc**—Frame Relay DLCI for a translational cross-connect **multilink-frame-relay-end-to-end**—IQE SONET PICs support Multilink Frame Relay end-to-end (FRF.15) **multilink-ppp**—IQE SONET PICs support Multilink PPP

Table 33: Encapsulation Support by Interface Type (continued)

Interface Type	Physical Interface Encapsulation	Logical Interface Encapsulation
t1—T1 interface (including channelized DS3-to-DS1 interfaces)	**cisco-hdlc**—Cisco-compatible HDLC framing **cisco-hdlc-ccc**—Cisco-compatible HDLC framing for a cross-connect **cisco-hdlc-tcc**—Cisco-compatible HDLC framing for a translational cross-connect **extended-frame-relay-ccc**—Any Frame Relay DLCI for a cross-connect **extended-frame-relay-tcc**—Any Frame Relay DLCI for a translational cross-connect **flexible-frame-relay**—Multiple Frame Relay encapsulations **frame-relay**—Frame Relay encapsulation **frame-relay-ccc**—Frame Relay for a cross-connect **frame-relay-port-ccc**—Frame Relay port encapsulation for a cross-connect **frame-relay-tcc**—Frame Relay for a translational cross-connect **multilink-frame-relay-uni-nni**—Multilink Frame Relay UNI NNI (FRF.16) encapsulation **ppp**—Serial PPP device **ppp-ccc**—Serial PPP device for a cross-connect **ppp-tcc**—Serial PPP device for a translational cross-connect	**frame-relay-ccc**—Frame Relay DLCI for CCC **frame-relay-ppp**—PPP over Frame Relay **frame-relay-tcc**—Frame Relay DLCI for a translational cross-connect

 Copyright © 2018, Juniper Networks, Inc.

Table 33: Encapsulation Support by Interface Type (continued)

Interface Type	Physical Interface Encapsulation	Logical Interface Encapsulation
t3—T3 interface (including channelized OC12-to-DS3 interfaces)	**cisco-hdlc**—Cisco-compatible HDLC framing **cisco-hdlc-ccc**—Cisco-compatible HDLC framing for a cross-connect **cisco-hdlc-tcc**—Cisco-compatible HDLC framing for a translational cross-connect **extended-frame-relay-ccc**—Any Frame Relay DLCI for a cross-connect **extended-frame-relay-tcc**—Any Frame Relay DLCI for a translational cross-connect **flexible-frame-relay**—Multiple Frame Relay encapsulations **frame-relay**—Frame Relay encapsulation **frame-relay-ccc**—Frame Relay for a cross-connect **frame-relay-port-ccc**—Frame Relay port encapsulation for a cross-connect **frame-relay-tcc**—Frame Relay for a translational cross-connect **ppp**—Serial PPP device **ppp-ccc**—Serial PPP device for a cross-connect **ppp-tcc**—Serial PPP device for a translational cross-connect	**frame-relay-ccc**—Frame Relay DLCI for CCC **frame-relay-ppp**—PPP over Frame Relay **frame-relay-tcc**—Frame Relay DLCI for a translational cross-connect
Controller-level channelized IQ interfaces (**cau4**, **coc1**, **coc3**, **coc12**, **cstm1**, **ct1**, **ct3**, **ce1**)	NA	NA
Services interfaces (**cp**, **gr**, **ip**, **mo**, **vt**, **es**, **mo**, **rsp**, **sp**)	NA	NA
Unconfigurable, internally generated interfaces (**gre**, **ipip**, **learning-chip (lc)**, **lsi**, **tap**, **mt**, **mtun**, **pd**, **pe**, **pimd**, **pime**)	NA	NA

NOTE: You can configure GRE interfaces (gre-x/y/z) only for GMPLS control channels. GRE interfaces are not supported or configurable for other applications. For more information about GMPLS, see the *Junos OS MPLS Applications Library for Routing Devices*.

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- Understanding Transient Interfaces on page 5
- Router Interfaces Overview on page 4
- Types of Interfaces Overview on page 4

Interface Descriptors Overview

When you configure an interface, you are effectively specifying the properties for a physical interface descriptor. In most cases, the physical interface descriptor corresponds to a single physical device and consists of the following parts:

- The interface name, which defines the media type
- The slot in which the FPC or DPC is located
- The location on the FPC in which the PIC is installed
- The PIC or DPC port
- The interface's channel and logical unit numbers (optional)

Each physical interface descriptor can contain one or more logical interface descriptors. These allow you to map one or more logical (or virtual) interfaces to a single physical device. Creating multiple logical interfaces is useful for ATM, Frame Relay, and Gigabit Ethernet networks, in which you can associate multiple virtual circuits, data-link connections, or virtual LANs (VLANs) with a single interface device.

Each logical interface descriptor can have one or more family descriptors to define the protocol family that is associated with and allowed to run over the logical interface.

The following protocol families are supported:

- Internet Protocol version 4 (IPv4) suite (inet)
- Internet Protocol version 6 (IPv6) suite (inet6)
- Circuit cross-connect (CCC)
- Translational cross-connect (TCC)
- International Organization for Standardization (ISO)
- Multilink Frame Relay end-to-end (MLFR end-to-end)
- Multilink Frame Relay user-to-network interface network-to-network interface (MLFR UNI NNI)
- Multilink Point-to-Point Protocol (MLPPP)
- Multiprotocol Label Switching (MPLS)
- Trivial Network Protocol (TNP)
- (M Series, T Series, and MX Series routers only) Virtual private LAN service (VPLS)

 Copyright © 2018, Juniper Networks, Inc.

Finally, each family descriptor can have one or more address entries, which associate a network address with a logical interface and hence with the physical interface.

You configure the various interface descriptors as follows:

- You configure the physical interface descriptor by including the **interfaces** ***interface-name*** statement.
- You configure the logical interface descriptor by including the **unit** statement within the **interfaces** ***interface-name*** statement or by including the ***.logical*** descriptor at the end of the interface name, as in **t3-0/0/0.1**, where the logical unit number is 1, as shown in the following examples:

```
[edit]
user@host# set interfaces t3-0/0/0 unit 1
[edit]
user@host# edit interfaces t3-0/0/0.1
[edit interfaces t3-0/0/0]
user@host# set unit 1
```

- You configure the family descriptor by including the **family** statement within the **unit** statement.
- You configure address entries by including the **address** statement within the **family** statement.
- You configure tunnels by including the **tunnel** statement within the **unit** statement.

NOTE: The address of a logical interface cannot be the same as a tunnel interface's source or destination address. If you try to configure a logical interface with a tunnel interface's address or vice versa, a commit failure will occur.

Related Documentation

- Router Interfaces Overview on page 4

Physical Part of an Interface Name

Interface Names for ACX Series Universal Metro Routers

ACX Series routers do not have actual PIC devices. Instead they have built-in network ports on the front panel of the router. These ports are named using the same naming convention used for routers with PIC devices with the understanding that the FPC, PIC and port are pseudo devices. When you display information about one of these ports, you specify the interface type, the slot for the Flexible PIC Concentrator (FPC), the slot on the FPC for the Physical Interface Card (PIC), and the configured port number.

In the physical part of the interface name, a hyphen (**-**) separates the media type from the FPC number, and a slash (**/**) separates the FPC, PIC, and port numbers:

type-fpc/pic/port

See Also

- *Understanding Encapsulation on an Interface*
- *Configuring Inverse Multiplexing for ATM (IMA) on ACX Series*

Interface Names for M Series and T Series Routers

On M Series and T Series routers, when you display information about an interface, you specify the interface type, the slot in which the Flexible PIC Concentrator (FPC) is installed, the slot on the FPC in which the Physical Interface Card (PIC) is located, and the configured port number.

In the physical part of the interface name, a hyphen (**-**) separates the media type from the FPC number, and a slash (**/**) separates the FPC, PIC, and port numbers:

type-fpc/pic/port

NOTE: **Exceptions to the *type-fpc/pic/port* physical description include the aggregated Ethernet and aggregated SONET/SDH interfaces, which use the syntax ae *number* and as *number*, respectively.**

MX Series Router Interface Names

On MX Series routers when you display information about an interface, you specify the interface type, the Dense Port Concentrator (DPC), Flexible PIC Concentrator (FPC), or Modular Port Concentrator (MPC) slot, the PIC or MIC slot, and the configured port number.

NOTE: **Although the MX Series routers use DPCs, FPCs, MPCs, MICs, and PICs, command syntax in this book is shown as *fpc/pic/port* for simplicity.**

In the physical part of the interface name, a hyphen (**-**) separates the media type from the FPC number, and a slash (**/**) separates the DPC, FPC or MPC, MIC or PIC, and port numbers:

type-fpc/pic/port

- *fpc*—Slot in which the DPC, FPC, or MPC is installed.
- *pic*—Slot on the FPC in which the PIC is located.

  For DPCs, MICs, and the 16-port MPC, the PIC value is a logical grouping of ports and varies on different platforms.

Copyright © 2018, Juniper Networks, Inc.

- *port*—Port number on the DPC, PIC, MPC, or MIC.

Interface Names for PTX Series Routers

On PTX Series Packet Transport Routers, when you display information about an interface, you specify the interface type, the slot in which the Flexible PIC Concentrator (FPC) is installed, the slot on the FPC in which the Physical Interface Card (PIC) is located, and the configured port number.

NOTE:

- The PTX router supports Ethernet type interfaces only. The media type portion of the physical interface name,*type* supports the Ethernet interface type only: **et**.
- In the CLI, all PTX3000 PICs are represented as **pic0**. For more information, see *PTX3000 Interface Modules*

In the physical part of the interface name, a hyphen (**-**) separates the media type (**et**) from the FPC number, and a slash (**/**) separates the FPC, PIC, and port numbers:

type-fpc/pic/port

Related Documentation

- Interface Naming Overview on page 32
- Logical Part of an Interface Name on page 39

Displaying Interface Configurations Overview

To display a configuration, use either the **show** command in configuration mode or the **show configuration** top-level command. Interfaces are listed in numerical order, from lowest to highest slot number, then from lowest to highest PIC number, and finally from lowest to highest port number.

Related Documentation

- Router Interfaces Overview on page 4

Copyright © 2018, Juniper Networks, Inc.

CHAPTER 2

Configuring Physical Interface Properties

Physical Interface Configuration Statements Overview

The software driver for each network media type sets reasonable default values for general interface properties, such as the interface's maximum transmission unit (MTU) size, receive and transmit leaky bucket properties, link operational mode, and clock source.

M Series, MX Series, and T Series routers are factory configured according to the specific router, its features, and its physical interfaces This section includes a default configuration example showing the statements used to configure the physical interfaces properties. Additional statements are used to set properties for specific interface types and are described in "Physical Interfaces Properties Statements List" on page 74.

To modify any of the default general interface properties, include the appropriate statements at the **[edit interfaces *interface-name*]** hierarchy level:

NOTE: The following configuration hierarchy and its included statements are shown only as an example of a configuration statement hierarchy and should not be referenced for resolving actual configurations. For information on a specific hierarchy level, see the hierarchy level document for that specific hierarchy, for example [edit interfaces] Hierarchy Level.

```
interfaces {
  traceoptions {
    file filename <files number> <match regular-expression> <size size> <world-readable |
      no-world-readable>;
    flag flag <disable>;
  }
  interface-name {
    accounting-profile name;
    aggregated-ether-options {
```

 Copyright © 2018, Juniper Networks, Inc.

```
    (flow-control | no-flow-control);
    lacp {
      (active | passive);
      link-protection{
        disable;
      (revertive | non-revertive (Interfaces));
      periodic interval;
      system-priority priority;
    }
    link-protection;
    link-speed speed;
    (loopback | no-loopback);
    minimum-links number;
    source-address-filter {
      mac-address
    }
    (source-filtering | no-source-filtering);
  }
  aggregated-sonet-options {
    link-speed speed | mixed;
    minimum-links number;
  }
  atm-options {
    cell-bundle-size cells;
    ilmi;
    linear-red-profiles profile-name {
      high-plp-max-threshold percent;
      low-plp-max-threshold percent;
      queue-depth cells high-plp-threshold percent low-plp-threshold percent;
    }
    mpls {
      pop-all-labels {
        required-depth number;
      }
    }
    pic-type (atm1 | atm2);
    plp-to-clp;
    promiscuous-mode {
      vpi vpi-identifier;
    }
    scheduler-maps map-name {
      forwarding-class class-name {
        epd-threshold cells plp1 cells;
        linear-red-profile profile-name;
        priority (high | low);
        transmit-weight (cells number | percent number);
      }
      vc-cos-mode (alternate | strict);
    }
    vpi vpi-identifier {
      maximum-vcs maximum-vcs;
      oam-liveness {
        up-count cells;
        down-count cells;
      }
      oam-period (seconds | disable);
```

```
    shaping {
      (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained
        rate burst length);
      queue-length number;
    }
  }
}
clocking clock-source;
data-input (system | interface interface-name);
dce;
serial-options {
  clock-rate rate;
  clocking-mode (dce | internal | loop);
  control-polarity (negative | positive);
  cts-polarity (negative | positive);
  dcd-polarity (negative | positive);
  dce-options {
    control-signal (assert | de-assert | normal);
    cts (ignore | normal | require);
    dcd (ignore | normal | require);
    dsr (ignore | normal | require);
    dtr signal-handling-option;
    ignore-all;
    indication (ignore | normal | require);
    rts (assert | de-assert | normal);
    tm (ignore | normal | require);
  }
  dsr-polarity (negative | positive);
  dte-options {
    control-signal (assert | de-assert | normal);
    cts (ignore | normal | require);
    dcd (ignore | normal | require);
    dsr (ignore | normal | require);
    dtr signal-handling-option;
    ignore-all;
    indication (ignore | normal | require);
    rts (assert | de-assert | normal);
    tm (ignore | normal | require);
  }
  dtr-circuit (balanced | unbalanced);
  dtr-polarity (negative | positive);
  encoding (nrz | nrzi);
  indication-polarity (negative | positive);
  line-protocol protocol;
  loopback mode;
  rts-polarity (negative | positive);
  tm-polarity (negative | positive);
  transmit-clock invert;
}
description text;
dialer-options {
  pool pool-name <priority priority>;
}
disable;
ds0-options {
  bert-algorithm algorithm;
```

 Copyright © 2018, Juniper Networks, Inc.

```
    bert-error-rate rate;
    bert-period seconds;
    byte-encoding (nx56 | nx64);
    fcs (16 | 32);
    idle-cycle-flag (flags | ones);
    invert-data;
    loopback payload;
    start-end-flag (filler | shared);
  }
  e1-options {
    bert-error-rate rate;
    bert-period seconds;
    fcs (16 | 32);
    framing (g704 | g704-no-crc4 | unframed);
    idle-cycle-flag (flags | ones);
    invert-data;
    loopback (local | remote);
    start-end-flag (filler | shared);
    timeslots time-slot-range;
  }
  e3-options {
    atm-encapsulation (direct | plcp);
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    buildout feet;
    compatibility-mode (digital-link | kentrox | larscom) <subrate value>;
    fcs (16 | 32);
    framing (g.751 | g.832);
    idle-cycle-flag (filler | shared);
    invert-data;
    loopback (local | remote);
    (payload-scrambler | no-payload-scrambler);
    start-end-flag (filler | shared);
    (unframed | no-unframed);
  }
  encapsulation type;
  es-options {
    backup-interface es-fpc/pic/port;
  }
  fastether-options {
    802.3ad aex;
    (flow-control | no-flow-control);
    ignore-l3-incompletes;
    ingress-rate-limit rate;
    (loopback | no-loopback);
    mpls {
      pop-all-labels {
        required-depth number;
      }
    }
    source-address-filter {
      mac-address;
    }
    (source-filtering | no-source-filtering);
  }
```

```
flexible-vlan-tagging;
gigether-options {
  802.3ad aex;
  (asynchronous-notification | no-asynchronous-notification);
  (auto-negotiation | no-auto-negotiation) remote-fault <local-interface-online |
    local-interface-offline>;
  auto-reconnect seconds;
  (flow-control | no-flow-control);
  ignore-l3-incompletes;
  (loopback | no-loopback);
  mpls {
    pop-all-labels {
      required-depth number;
    }
  }
  source-address-filter {
    mac-address;
  }
  (source-filtering | no-source-filtering);
  ethernet-switch-profile {
    (mac-learn-enable | no-mac-learn-enable);
    tag-protocol-id [ tpids ];
    ethernet-policer-profile {
      input-priority-map {
        ieee802.1p premium [ values ];
      }
      output-priority-map {
        classifier {
          premium {
            forwarding-class class-name {
              loss-priority (high | low);
            }
          }
        }
      }
      policer cos-policer-name {
        aggregate {
          bandwidth-limit bps;
          burst-size-limit bytes;
        }
        premium {
          bandwidth-limit bps;
          burst-size-limit bytes;
        }
      }
    }
  }
}
(gratuitous-arp-reply | no-gratuitous-arp-reply);
hold-time up milliseconds down milliseconds;
interface-set interface-set-name {
  interface ethernet-interface-name {
    (unit unit-number | vlan-tags-outer vlan-tag);
  }
}
isdn-options {
```

 Copyright © 2018, Juniper Networks, Inc.

```
    bchannel-allocation (ascending | descending);
    calling-number number;
    pool pool-name <priority priority>;
    spid1 spid-string;
    spid2 spid-string;
    static-tei-val value;
    switch-type (att5e | etsi | ni1 | ntdms100 | ntt);
    t310 seconds;
    tei-option (first-call | power-up);
  }
  keepalives <down-count number> <interval seconds> <up-count number>;
  link-mode mode;
  lmi {
    lmi-type (ansi | itu | c-lmi);
    n391dte number;
    n392dce number;
    n392dte number;
    n393dce number;
    n393dte number;
    t391dte seconds;
    t392dce seconds;
  }
  lsq-failure-options {
    no-termination-request;
    [ trigger-link-failure interface-name ];
  }
  mac mac-address;
  mlfr-uni-nni-bundle-options {
    acknowledge-retries number;
    acknowledge-timer milliseconds;
    action-red-differential-delay (disable-tx | remove-link);
    cisco-interoperability send-lip-remove-link-for-link-reject;
    drop-timeout milliseconds;
    fragment-threshold bytes;
    hello-timer milliseconds;
    link-layer-overhead percent;
    lmi-type (ansi | itu | c-lmi);
    minimum-links number;
    mrru bytes;
    n391 number;
    n392 number;
    n393 number;
    red-differential-delay milliseconds;
    t391 seconds;
    t392 seconds;
    yellow-differential-delay milliseconds;
    encapsulation type;
  }
  modem-options {
    dialin (console | routable);
    init-command-string initialization-command-string;
  }
  mtu bytes;
  multiservice-options {
    (core-dump | no-core-dump);
    (syslog | no-syslog);
```

Copyright © 2018, Juniper Networks, Inc.

```
    (dump-on-flow-control);
    flow-control-options {
      down-on-flow-control;
      dump-on-flow-control;
      reset-on-flow-control;
    }
  }
  native-vlan-id number;
  no-gratuitous-arp-request;
  no-keepalives;
  no-partition {
    interface-type type;
  }
  optics-options {
    wavelength nm;
    alarmalarm-name {
      (syslog | link-down);
    }
    warningwarning-name {
      (syslog | link-down);
    }
  }
  partition partition-number oc-slice oc-slice-range interface-type type;
  timeslots time-slot-range;
  passive-monitor-mode;
  per-unit-scheduler;
  ppp-options {
    chap {
      access-profile name;
      default-chap-secret name;
      local-name name;
      passive;
    }
    compression {
      acfc;
      pfc;
    }
    dynamic-profile profile-name;
    no-termination-request;
    pap {
      access-profile name;
      local-name name;
      local-password password;
      passive;
    }
  }
  receive-bucket {
    overflow (discard | tag);
    rate percentage;
    threshold bytes;
  }
  redundancy-options {
    primary sp-fpc/pic/port;
    secondary sp-fpc/pic/port;
  }
  schedulers number;
```

 Copyright © 2018, Juniper Networks, Inc.

```
serial-options {
  clock-rate rate;
  clocking-mode (dce | internal | loop);
  control-polarity (negative | positive);
  cts-polarity (negative | positive);
  dcd-polarity (negative | positive);
  dce-options {
    control-signal (assert | de-assert | normal);
    cts (ignore | normal | require);
    dcd (ignore | normal | require);
    dsr (ignore | normal | require);
    dtr signal-handling-option;
    ignore-all;
    indication (ignore | normal | require);
    rts (assert | de-assert | normal);
    tm (ignore | normal | require);
  }
  dsr-polarity (negative | positive);
  dte-options {
    control-signal (assert | de-assert | normal);
    cts (ignore | normal | require);
    dcd (ignore | normal | require);
    dsr (ignore | normal | require);
    dtr signal-handling-option;
    ignore-all;
    indication (ignore | normal | require);
    rts (assert | de-assert | normal);
    tm (ignore | normal | require);
  }
  dtr-circuit (balanced | unbalanced);
  dtr-polarity (negative | positive);
  encoding (nrz | nrzi);
  indication-polarity (negative | positive);
  line-protocol protocol;
  loopbackmode;
  rts-polarity (negative | positive);
  tm-polarity (negative | positive);
  transmit-clock invert;
}
services-options {
  inactivity-timeout seconds;
  open-timeout seconds;
  syslog {
    host hostname {
      facility-override facility-name;
      log-prefix prefix-number;
      services priority-level;
    }
  }
}
shdsl-options {
  annex (annex-a | annex-b);
  line-rate line-rate;
  loopback (local | remote);
  snr-margin {
    snext margin;
```

```
    }
  }
  sonet-options {
    aggregate asx;
    aps {
      advertise-interval milliseconds;
      authentication-key key;
      force;
      hold-time milliseconds;
      lockout;
      neighbor address;
      paired-group group-name;
      preserve-interface;
      protect-circuit group-name;
      request;
      revert-time seconds;
      switching-mode (bidirectional | unidirectional);
      working-circuit group-name;
    }
    bytes {
      c2 value;
      e1-quiet value;
      f1 value;
      f2 value;
      s1 value;
      z3 value;
      z4 value;
    }
    fcs (16 | 32);
    loopback (local | remote);
    mpls {
      pop-all-labels {
        required-depth number;
      }
    }
    path-trace trace-string;
    (payload-scrambler | no-payload-scrambler);
    rfc-2615;
    trigger {
      defect ignore;
      hold-time up milliseconds down milliseconds;
    }
    vtmapping (itu-t | klm);
    (z0-increment | no-z0-increment);
  }
  (speed (10m | 100m | 1g | auto) | speed (auto | 1Gbps | 100Mbps | 10Mbps) | speed
    (oc3 | oc12 | oc48));
  stacked-vlan-tagging;
  switch-options {
    switch-port port-number {
      (auto-negotiation | no-auto-negotiation);
      speed (10m | 100m | 1g);
      link-mode (full-duplex | half-duplex);
    }
  }
  multicast-statistics
```

 Copyright © 2018, Juniper Networks, Inc.

```
t1-options {
  bert-algorithm algorithm;
  bert-error-rate rate;
  bert-period seconds;
  buildout value;
  byte-encoding (nx56 | nx64);
  crc-major-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5);
  crc-minor-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5 | 5e-6 | 1e-6);
  fcs (16 | 32);
  framing (esf | sf);
  idle-cycle-flag (flags | ones);
  invert-data;
  line-encoding (ami | b8zs);
  loopback (local | payload | remote);
  remote-loopback-respond;
  start-end-flag (filler | shared);
  timeslots time-slot-range;
}
t3-options {
  atm-encapsulation (direct | plcp);
  bert-algorithm algorithm;
  bert-error-rate rate;
  bert-period seconds;
  buildout feet;
  (cbit-parity | no-cbit-parity);
  compatibility-mode (adtran | digital-link | kentrox | larscom | verilink) <subrate
    value>;
  fcs (16 | 32);
  (feac-loop-respond | no-feac-loop-respond);
  idle-cycle-flag value;
  (long-buildout | no-long-buildout);
  (loop-timing | no-loop-timing);
  loopback (local | payload | remote);
  (mac | no-mac);
  (payload-scrambler | no-payload-scrambler);
  start-end-flag (filler | shared);
}
traceoptions {
  flag flag <flag-modifier> <disable>;
}
transmit-bucket {
  overflow discard;
  rate percentage;
  threshold bytes;
}
(traps | no-traps);
unidirectional;
vlan-tagging;
vlan-vci-tagging;
unit logical-unit-number {
  logical-interface-statements;
}
```
```
  }
}
```

Related Documentation

- Router Interfaces Overview on page 4

Physical Interfaces Properties Statements List

Table 34 on page 74 lists statements that you can use to configure physical interfaces.

Table 34: Statements for Physical Interface Properties

Statement	Interface Types	Usage Guidelines
802.3ad aex	Aggregated Ethernet interfaces	*Aggregated Ethernet Interfaces Overview*
access-profile *name*	Interfaces with Point-to-Point Protocol (PPP) encapsulation	"Configuring the PPP Challenge Handshake Authentication Protocol" on page 136
accounting-profile *name*	All	"Configuring Accounting for the Physical Interface" on page 171
acfc	Interfaces with PPP encapsulation	*Identifying the Access Concentrator*
acknowledge-retries *number*	Link services and voice services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
acknowledge-timer *milliseconds*	Link services and voice services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
action-red-differential-delay (disable-tx | remove-link)	Link services and voice services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
advertise-interval *milliseconds*	SONET/SDH interfaces	*Configuring APS Timers*
aggregate	Gigabit Ethernet intelligent queuing (IQ and IQE) interfaces and Gigabit Ethernet interfaces with small form-factor pluggable transceivers (SFPs) (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router)	*Configuring Gigabit Ethernet Policers*
aggregate as*x*	Aggregated SONET/SDH interfaces	*Configuring Aggregated SONET/SDH Interfaces*
aggregated-ether-options	Aggregated Ethernet interfaces	*Aggregated Ethernet Interfaces Overview*
aggregate-ports	SONET/SDH interfaces	*Configuring 4-Port OC192 PIC to Operate in OC768-over-OC192 Mode*
aggregated-sonet-options	Aggregated SONET/SDH interfaces	*Configuring Aggregated SONET/SDH Interfaces*

 Copyright © 2018, Juniper Networks, Inc.

Table 34: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
alarm*alarm-name* *(syslog | link-down)*	10-Gigabit Ethernet interfaces	*Configuring 10-Gigabit Ethernet Link Down Notification for Optics Options Alarm or Warning*
alias *alias-name*;	All	"Example: Adding an Interface Alias Name" on page 123
annex (annex-a | annex-b)	SONET interfaces using annex-b for MSP switching on M320 and M120 Routers	*Configuring Basic Automatic Protect Switching*
aps	SONET/SDH interfaces	*SONET/SDH Interfaces Overview*
atm-encapsulation (direct | plcp)	E3 and T3 traffic over Asynchronous Transfer Mode (ATM) interfaces	*Configuring E3 and T3 Parameters on ATM Interfaces*
atm-options	ATM1 and ATM2 IQ interfaces	"Interface Encapsulations Overview" on page 46
authentication-key *key*	SONET/SDH interfaces	*SONET/SDH Interfaces Overview*
bandwidth-limit *bps*	Gigabit Ethernet and Gigabit Ethernet IQ and IQE PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router)	*Configuring Gigabit Ethernet Policers*
bert-algorithm *algorithm*	E3, T1, T3, multichannel DS3, channelized interfaces (DS3, OC12, and STM1), and channelized IQ and IQE interfaces (E1 and DS3)	*Configuring Interface Diagnostics Tools to Test the Physical Layer Connections*
bert-error-rate *rate*	E1, E3, T1, T3, and channelized interfaces (DS3, OC3, OC12, and STM1)	*Configuring Interface Diagnostics Tools to Test the Physical Layer Connections*
bert-period *seconds*	E1, E3, T1, T3, and channelized interfaces (DS3, OC12, and STM1)	*Configuring Interface Diagnostics Tools to Test the Physical Layer Connections*
Configuring the T1 Buildout *value*	T1 interfaces	*Configuring the T1 Buildout*
buildout *feet*	E3 and T3 traffic over ATM interfaces	*Configuring E3 and T3 Parameters on ATM Interfaces*
burst-size-limit *bytes*	Gigabit Ethernet and Gigabit Ethernet IQ and IQE PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router)	*Configuring Gigabit Ethernet Policers*
byte-encoding (nx56 | nx64)	DS0 and T1 interfaces	*Configuring T1 Byte Encoding*

Copyright © 2018, Juniper Networks, Inc.

Table 34: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
bytes [*values*	SONET/SDH interfaces	*Configuring SONET/SDH Header Byte Values to Identify Error Conditions*
cbit-parity | no-cbit-parity	T3 interfaces	*Disabling T3 C-Bit Parity Mode*
cbr *rate*	ATM interfaces	*Defining the ATM Traffic-Shaping Profile Overview*
cell-bundle-size *cells*	ATM2 IQ interfaces using ATM Layer 2 circuit cell-relay transport mode	*Configuring the Layer 2 Circuit Cell-Relay Cell Maximum Overview*
chap	Interfaces with PPP encapsulation	"Configuring the PPP Challenge Handshake Authentication Protocol" on page 136
cisco-interoperability send-lip-remove-link-for-link-reject	link services IQ (lsq) interfaces	*Junos OS Services Interfaces Library for Routing Devices*
classifier	Gigabit Ethernet IQ interfaces	*Configuring Gigabit Ethernet Policers*
clocking *clock-source*	ATM, DS0, E1, E3, SONET/SDH, T1, and T3 interfaces	"Configuring the Clock Source" on page 128
clocking-mode (dce | internal | loop)	Serial interfaces (EIA-530 and V.35)	"Configuring the Serial Clocking Mode" on page 342
clock-rate *rate*	Serial interfaces (EIA-530 and V.35)	"Configuring the DTE Clock Rate" on page 343
compatibility-mode *mode*	E3 and T3 interfaces	*Configuring the E3 CSU Compatibility Mode* and *Configuring the T3 CSU Compatibility Mode*
compression	Interfaces with PPP encapsulation	"Configuring the PPP Protocol Field Compression" on page 151
control-polarity (negative | positive)	Serial interfaces (X.21)	"Configuring Serial Signal Polarities" on page 347
control-signal (assert | de-assert | normal)	Serial interfaces (X.21)	"Configuring the Serial Signal Handling" on page 344
core-dump | no-core-dump)	Adaptive services, monitoring services, and collector interfaces	"Configuring Multiservice Physical Interface Properties" on page 170
cts (ignore | normal | require)	Serial interfaces (EIA-530 and V.35)	"Configuring the Serial Signal Handling" on page 344
cts-polarity (negative | positive)	Serial interfaces (EIA-530 and V.35)	"Configuring Serial Signal Polarities" on page 347

 Copyright © 2018, Juniper Networks, Inc.

Table 34: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
dcd (ignore | normal | require)	Serial interfaces (EIA-530 and V.35)	"Configuring the Serial Signal Handling" on page 344
dcd-polarity (negative | positive)	Serial interfaces (EIA-530 and V.35)	"Configuring Serial Signal Polarities" on page 347
dce	Interfaces with Frame Relay encapsulation	"Configuring the Router as a DCE with Frame Relay Encapsulation" on page 154
default-chap-secret *name*	Interfaces with Point-to-Point Protocol (PPP) encapsulation	*Configuring the PPP Challenge Handshake Authentication Protocol*
description *text*	All	"Configuring Interface Description" on page 113
disable	All	"Disabling a Physical Interface" on page 174 and "Tracing Operations of an Individual Router Interface" on page 355
dot1x	802.1x Port-Based Network Access Control	*IEEE 802.1x Port-Based Network Access Control Overview*
down-count	ATM interfaces	*Configuring the ATM OAM F5 Loopback Cell Threshold*
drop-timeout *milliseconds*	Multilink, link services, and voice services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
ds0-options	DS0 interfaces	*Channelized Interfaces Overview*
dsr (ignore | normal | require)	Serial interfaces (EIA-530 and V.35)	"Configuring the Serial Signal Handling" on page 344
dsr-polarity (negative | positive)	Serial interfaces (EIA-530 and V.35)	"Configuring Serial Signal Polarities" on page 347
dte-options	Serial interfaces (EIA-530, V.35, and X.21) on M Series and T Series routers	"Configuring the Serial Signal Handling" on page 344
dtr *signal-handling-option*	Serial interfaces (EIA-530 and V.35)	"Configuring the Serial Signal Handling" on page 344
dtr-circuit (balanced | unbalanced)	Serial interfaces (EIA-530 and V.35)	"Configuring the Serial DTR Circuit" on page 347
dtr-polarity (negative | positive)	Serial interfaces (EIA-530 and V.35)	"Configuring Serial Signal Polarities" on page 347

Copyright © 2018, Juniper Networks, Inc.

Table 34: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
e1-options	E1 interfaces	*E1 Interfaces Overview*
e3-options	E3 interfaces	*E3 Interfaces Overview*
encapsulation *type*	All interfaces, except loopback and multicast tunnel	"Configuring Interface Encapsulation on Physical Interfaces" on page 129
encoding (nrz | nrzi)	Serial interfaces (EIA-530, V.35, and X.21)	"Configuring Serial Line Encoding" on page 350
epd-threshold *cells*	ATM2 interfaces	*ATM2 IQ VC Tunnel CoS Components Overview*
es-options	ES interfaces	*Junos OS Services Interfaces Library for Routing Devices*
ethernet-policer-profile	Gigabit Ethernet and Gigabit Ethernet IQ and IQE PICs with SFPs (except the 10-port Gigabit Ethernet PIC, and the built-in Gigabit Ethernet port on the M7i router)	*Configuring Gigabit Ethernet Policers*
ethernet-switch-profile	Gigabit Ethernet and Gigabit Ethernet IQ and IQE PICs with SFPs (except the 10-port Gigabit Ethernet PIC, Aggregated Ethernet with Gigabit Ethernet IQ interfaces, and the built-in Gigabit Ethernet port on the M7i router)	*Configuring Gigabit Ethernet Policers*, *Configuring MAC Address Filtering*, and Configuring the Management Ethernet Interface
facility-override *facility-name*	Adaptive services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
fastether-options	Fast Ethernet interfaces	*Ethernet Interfaces Overview*
fcs (16 | 32)	E1/E3, SONET/SDH, and T1/T3 interfaces	*Configuring the E1 Frame Checksum*, *Configuring the E3 Frame Checksum*, *Configuring the SONET/SDH Frame Checksum*, *Configuring the T1 Frame Checksum*, and *Configuring the T3 Frame Checksum*
feac-loop-respond | no-feac-loop-respond)	T3 interfaces	*Configuring the T3 FEAC Response*
flow-control | no-flow-control)	Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces	*Configuring Flow Control*
force	SONET/SDH interfaces	*Configuring Switching Between the Working and Protect Circuits*

Copyright © 2018, Juniper Networks, Inc.

Table 34: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
forwarding-class *class-name*	ATM2 IQ interfaces	*ATM2 IQ VC Tunnel CoS Components Overview*
forwarding-class *class-name*	Gigabit Ethernet IQ interfaces	*Configuring Gigabit Ethernet Policers*
fragment-threshold *bytes*	Multilink, link services, and voice services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
framing*framing-type*	E1, E3, and T1 interfaces	*Configuring E3 and T3 Parameters on ATM Interfaces*, *Configuring E1 Framing*, and *Configuring T1 Framing*
framing *framing-type*	10-Gigabit Ethernet interfaces	*10-Gigabit Ethernet Framing Overview*
framing *framing-type*	SONET interfaces	*Configuring SONET/SDH Framing Mode for Ports*
gigether-options	Gigabit Ethernet and Tri-Rate Ethernet copper interfaces	*Ethernet Interfaces Overview*
(gratuitous-arp-reply | no-gratuitous-arp-reply)	Ethernet interfaces	*Configuring Gratuitous ARP*
hello-timer *milliseconds*	Link services and voice services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
high-plp-max-threshold	ATM2 interfaces	*Configuring ATM2 IQ VC Tunnel CoS Components*
high-plp-threshold *percent*	ATM2 interfaces	*ATM2 IQ VC Tunnel CoS Components Overview*
hold-time *milliseconds*	SONET/SDH interfaces	*SONET/SDH Interfaces Overview*
hold-time up *milliseconds* down *milliseconds*	All interfaces, except aggregated SONET/SDH, generalized routing encapsulation (GRE) tunnel, and IP tunnel	*Configuring SONET/SDH Defect Triggers*
host *hostname*	Adaptive services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
ieee802.1p premium [*values*]	Gigabit Ethernet IQ interfaces	*Specifying an Input Priority Map*
idle-cycle-flag *value*	E1, E3, T1, and T3 interfaces	*Configuring the E1 Idle Cycle Flag*, *Configuring the E3 Idle Cycle Flag*, *Configuring the T1 Idle Cycle Flag*, and *Configuring the T3 Idle Cycle Flag*
ignore-all	Serial interfaces (EIA-530, V.35, and X.21)	"Configuring the Serial Signal Handling" on page 344

Copyright © 2018, Juniper Networks, Inc.

Table 34: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
ilmi	ATM interfaces	*Configuring Communication with Directly Attached ATM Switches and Routers*
inactivity-timeout *seconds*	Adaptive services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
indication (ignore | normal | require)	Serial interfaces (X.21)	"Configuring the Serial Signal Handling" on page 344
indication-polarity (negative | positive)	Serial interfaces (X.21)	"Configuring Serial Signal Polarities" on page 347
ingress-rate-limit *rate*	8-port, 12-port, and 48-port Fast Ethernet interfaces	*Configuring the Ingress Rate Limit*
input-priority-map	Gigabit Ethernet IQ interfaces	*Specifying an Input Priority Map*
interface-type *type*	Channelized IQ and IQE interfaces	*Channelized Interfaces Overview*
invert-data	DS0, E1, E3, and T1 interfaces	*Configuring E1 Data Inversion*, *Configuring E3 Data Inversion*, and *Configuring T1 Data Inversion*
keepalives <down-count *number* <interval *seconds*> <up-count *number*>	Aggregated SONET/SDH, DS0, E1, E3, SONET/SDH, T1, and T3 interfaces	"Configuring Keepalives" on page 134
lacp *mode*	Aggregated Ethernet interfaces	*Configuring LACP for Aggregated Ethernet Interfaces*
line-encoding (ami | b8zs)	T1 interfaces	*Configuring T1 Line Encoding*
line-protocol *protocol*	Serial interfaces (EIA-530, V.35, and X.21)	"Configuring the Serial Line Protocol" on page 338
linear-red-profile *profile-name*	ATM2 interfaces	*ATM2 IQ VC Tunnel CoS Components Overview*
linear-red-profiles *profile-name*	ATM2 interfaces	*ATM2 IQ VC Tunnel CoS Components Overview*
link-layer-overhead *percent*	AS PIC link services IQ interfaces (**lsq**)	*Junos OS Services Interfaces Library for Routing Devices*
link-mode *mode*	Management Ethernet (**fxp0** or **em0**) and Fast Ethernet interfaces	"Configuring the Link Characteristics" on page 121, "Understanding Management Ethernet Interfaces" on page 10

Copyright © 2018, Juniper Networks, Inc.

Table 34: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
link-speed *speed*	Aggregated Ethernet interfaces	"Configuring Aggregated Ethernet Link Speed" on page 118
link-speed *speed*	Aggregated SONET/SDH interfaces	*Configuring Aggregated SONET/SDH Interfaces*
lmi *lmi-options*	Interfaces with Frame Relay encapsulation	*Configuring Frame Relay Keepalives* and *Junos OS Services Interfaces Library for Routing Devices*
lmi	OAM CFM Ethernet Local Management Interface	*Configuring Ethernet Local Management Interface*
lmi-type (ansi | itu | c-lmi)	Link services interfaces and interfaces with Frame Relay encapsulation	*Configuring Frame Relay Keepalives*
local-name *name*	Interfaces with PPP encapsulation	"Configuring the PPP Challenge Handshake Authentication Protocol" on page 136
lockout	SONET/SDH interfaces	*Configuring Lockout of Protection for SDH Interfaces*
log-prefix *prefix-number*	Adaptive services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
(long-buildout | no-long-buildout)	T3 interfaces	*Configuring the T3 Line Buildout*
(loop-timing | no-loop-timing)	Channelized IQ interfaces	*Configuring the Channelized T3 Loop Timing*
loopback *mode*	DS0, E1, E3, T1, T3, SHDSL on ATM and SONET/SDH interfaces	*Configuring E1 Loopback Capability*, *Configuring E3 Loopback Capability*, *Configuring T1 Loopback Capability*, *Configuring T3 Loopback Capability*, *Configuring SHDSL Operating Mode on an ATM Physical Interface*,*Configuring SONET/SDH Loopback Capability to Identify a Problem as Internal or External*, and *Configuring Ethernet Loopback Capability*
loopback *mode*	Ethernet and 10-Gigabit Ethernet interfaces in WAN PHY mode	*Configuring Ethernet Loopback Capability*
loopback *mode*	Serial interfaces	"Configuring Serial Loopback Capability" on page 348
(loopback | no-loopback)	Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces	*Configuring Ethernet Loopback Capability*
loss-priority (high | low)	Gigabit Ethernet IQ interfaces	*Configuring Gigabit Ethernet Policers*

Copyright © 2018, Juniper Networks, Inc.

Table 34: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
lowest-priority-defect (all-defects | err-xcon | mac-rem-err-xcon | no-defect | rem-err-xcon | xcon)	Configuring IEEE 802.1ag OAM connectivity-fault management	*Configuring a MEP to Generate and Respond to CFM Protocol Messages*
low-plp-max-threshold *percent*	ATM2 interfaces	*ATM2 IQ VC Tunnel CoS Components Overview*
low-plp-threshold *percent*	ATM2 interfaces	*ATM2 IQ VC Tunnel CoS Components Overview*
lsq-failure-options	Link services IQ (lsq) interfaces	*Junos OS Services Interfaces Library for Routing Devices*
mac *mac-address*	Management Ethernet interface (**fxp0** or **em0**)	*Configuring the MAC Address on the Management Ethernet Interface*, "Understanding Management Ethernet Interfaces" on page 10
(mac-learn-enable | no-mac-learn-enable)	Gigabit Ethernet IQ and IQE, Tri-Rate Ethernet copper, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router)	*Configuring MAC Address Filtering*
master-only;	Management Ethernet (**fxp0** or **em0**) and Fast Ethernet interfaces	*Configuring a Consistent Management IP Address*, "Understanding Management Ethernet Interfaces" on page 10
maximum-vcs *maximum-vcs*	ATM interfaces	*Configuring the Maximum Number of ATM1 VCs on a VP*
mc-ae	Aggregated Ethernet interfaces	*Configuring Multichassis Link Aggregation on MX Series Routers*
minimum-links *number*	Multilink, link services, and voice services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
mip-half-function	Connectivity Fault Management	*Configuring IEEE 802.3ah OAM Link-Fault Management*
mlfr-uni-nni-bundle-options *bundle-options*	Link services and voice services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
mpls	10-Gigabit Ethernet interfaces in WAN PHY mode and ATM and SONET/SDH interfaces in passive monitoring mode	*Removing MPLS Labels from Incoming Packets*, *Enabling Packet Flow Monitoring on SONET/SDH Interfaces*, and *SONET/SDH Interfaces Overview*
mrru *bytes*	Link services and voice services interfaces	*Junos Services Interfaces Configuration Guide*

Copyright © 2018, Juniper Networks, Inc.

Table 34: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
mtu *bytes*	All interfaces, except management Ethernet (**fxp0** or **em0**), loopback, multilink, and multicast tunnel	"Understanding Management Ethernet Interfaces" on page 10, "Media MTU Overview" on page 100, "Configuring the Media MTU" on page 115
multicast-statistics	Ethernet, SONET, aggregated Ethernet, and aggregated SONET interfaces.	*Configuring Multicast Statistics Collection on Aggregated Ethernet Interfaces*, *Configuring Multicast Statistics Collection on SONET Interfaces*,*Configuring Multicast Statistics Collection on Ethernet Interfaces*, and *Configuring Multicast Statistics Collection on Aggregated SONET Interfaces*
multiservice-options	Adaptive services, monitoring services, and collector interfaces	"Configuring Multiservice Physical Interface Properties" on page 170
n391 *number*	Link services and voice services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
n392 *number*	Link services and voice services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
n393 *number*	Link services and voice services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
neighbor *address*	SONET/SDH interfaces	*SONET/SDH Interfaces Overview*
no-gratuitous-arp-request	Ethernet interfaces	*Configuring Gratuitous ARP*
no-keepalives	Interfaces with PPP, Frame Relay, or Cisco High-level Data Link Control (HDLC) encapsulation	"Configuring Keepalives" on page 134
no-partition	Channelized IQ interfaces	*Channelized Interfaces Overview*
no-termination-request	Link Services IQ (LSQ) interfaces	*Junos OS Services Interfaces Library for Routing Devices*
oam-liveness	ATM interfaces	*Configuring the OAM F4 Cell Flows*
oam-period (*seconds* | disable)	ATM interfaces	*Defining the ATM OAM F5 Loopback Cell Period*
oc-slice *oc-slice-range*	Channelized OC12 IQ interfaces	*Channelized OC12/STM4 IQ and IQE Interfaces Overview*
open-timeout *seconds*	Adaptive services interfaces	*Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2018, Juniper Networks, Inc.

Table 34: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
optics-options	Gigabit Ethernet dense wavelength-division multiplexing (DWDM) interfaces	*Ethernet DWDM Interface Wavelength Overview*
output-priority-map	Gigabit Ethernet IQ interfaces	*Configuring Gigabit Ethernet Policers*
overflow (discard | tag) (Receive bucket)	All interfaces, except ATM, channelized E1, E1, Fast Ethernet, Gigabit Ethernet, and channelized IQ interfaces	"Configuring Receive and Transmit Leaky Bucket Properties to Reduce Network Congestion" on page 156
overflow (discard) (Transmit bucket)		
paired-group *group-name*	SONET/SDH interfaces	*Configuring APS Load Sharing*
partition *partition-number*	Channelized IQ interfaces	*Channelized Interfaces Overview*
passive	Interfaces with PPP encapsulation	"Configuring the PPP Challenge Handshake Authentication Protocol" on page 136
passive-monitor-mode	SONET/SDH interfaces	*Enabling Packet Flow Monitoring on SONET/SDH Interfaces*
path-trace *trace-string*	10-Gigabit Ethernet interfaces in WAN PHY mode and SONET/SDH interfaces	*Configuring the SONET/SDH Path Trace Identifier for a Circuit*
(payload-scrambler | no-payload-scrambler)	E3, SONET/SDH, and T3 interfaces	*Configuring E3 and T3 Parameters on ATM Interfaces*, *Configuring E3 HDLC Payload Scrambling*, *Configuring SONET/SDH HDLC Payload Scrambling for Link Stability*, and *Configuring T3 HDLC Payload Scrambling*
periodic *interval*	Aggregated Ethernet interfaces	*Configuring LACP for Aggregated Ethernet Interfaces*
per-unit-scheduler	IQ interfaces	*Junos OS Services Interfaces Library for Routing Devices*
pfc	Interfaces with PPP encapsulation	"Configuring the PPP Protocol Field Compression" on page 151
pic-type (atm1 | atm2)	ATM2 IQ interfaces	*Configuring the ATM PIC Type*
plp1 *cells*	ATM2 interfaces	*ATM2 IQ VC Tunnel CoS Components Overview*
plp-to-clp	ATM2 IQ interfaces	*Enabling the PLP Setting to Be Copied to the CLP Bit*

Copyright © 2018, Juniper Networks, Inc.

Table 34: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
policer *cos-policer-name*	Gigabit Ethernet and Gigabit Ethernet IQ and IQE PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router)	*Configuring Gigabit Ethernet Policers*
pop-all-labels	ATM and SONET/SDH interfaces in passive monitoring mode	*Removing MPLS Labels from Incoming Packets* and *Enabling Packet Flow Monitoring on SONET/SDH Interfaces*
ppp-options	Interfaces with PPP encapsulation	"Configuring the PPP Challenge Handshake Authentication Protocol" on page 136
premium	Enhanced Intelligent Queuing (IQE) interfaces (hierarchical policer)	"Applying Policers" on page 226 and *Class of Service Feature Guide for Routing Devices and EX9200 Switches*
premium	Gigabit Ethernet IQ interfaces (policer)	*Configuring Gigabit Ethernet Policers*
premium	Gigabit Ethernet IQ interfaces (output priority map)	*Configuring MAC Address Filtering*
primary sp-*fpc/pic/port*	Redundant interfaces for adaptive services interfaces (**rsp-**)	*Junos OS Services Interfaces Library for Routing Devices*
priority (high | low)	ATM2 IQ interfaces	*ATM2 IQ VC Tunnel CoS Components Overview*
priority *number*	Ethernet protocols (OAM CFM)	*Configuring a MEP to Generate and Respond to CFM Protocol Messages*
promiscuous-mode	ATM2 IQ interfaces	*Configuring ATM Cell-Relay Promiscuous Mode*
protect-circuit *group-name*	SONET/SDH interfaces	*Configuring Switching Between the Working and Protect Circuits*
queue-depth *cells*	ATM2 interfaces	*ATM2 IQ VC Tunnel CoS Components Overview*
queue-length *number*	ATM1 interfaces	*Configuring the ATM1 Queue Length*
rate *percentage*	All interfaces, except ATM, channelized E1, E1, Fast Ethernet, Gigabit Ethernet, and channelized IQ	"Configuring Receive and Transmit Leaky Bucket Properties to Reduce Network Congestion" on page 156
receive-bucket	All interfaces, except ATM, Fast Ethernet, and Gigabit Ethernet	"Configuring Receive and Transmit Leaky Bucket Properties to Reduce Network Congestion" on page 156

Copyright © 2018, Juniper Networks, Inc.

Table 34: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
red-differential-delay *milliseconds*	Link services and voice services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
redundancy-options	Redundant interfaces for adaptive services interfaces (**rsp-**)	*Junos OS Services Interfaces Library for Routing Devices*
remote-loopback-respond	T1 interfaces	*Configuring the T1 Remote Loopback Response*
request	SONET/SDH interfaces	*Configuring Switching Between the Working and Protect Circuits*
required-depth *number*	ATM and SONET/SDH interfaces in passive monitoring mode	*Removing MPLS Labels from Incoming Packets* and *Enabling Packet Flow Monitoring on SONET/SDH Interfaces*
revert-time *seconds*	SONET/SDH interfaces	*Configuring Revertive Mode*
rfc-2615	SONET/SDH interfaces	*Configuring PPP Support on SONET/SDH Interfaces*
rts (assert | de-assert | normal)	Serial interfaces (EIA-530 and V.35)	"Configuring the Serial Signal Handling" on page 344
rts-polarity (negative | positive)	Serial interfaces (EIA-530 and V.35)	"Configuring Serial Signal Polarities" on page 347
rtvbr peak *rate* sustained *rate* burst *length*	ATM interfaces	*Configuring ATM CBR*
scheduler-maps *map-name*	ATM2 interfaces	*ATM2 IQ VC Tunnel CoS Components Overview*
schedulers *number*	Ethernet IQ2 and IQ2-E PICs port interfaces	*Junos OS Services Interfaces Library for Routing Devices*
secondary sp-*fpc/pic/port*	Redundant interfaces for adaptive services interfaces (**rsp-**)	*Junos OS Services Interfaces Library for Routing Devices*
services-options	Services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
serial-options	Serial interfaces (EIA-530, V.35, and X.21)	"Serial Interfaces Overview" on page 335
services *priority-level*	Adaptive services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
size	All	"Tracing Operations of the Interface Process" on page 356

 Copyright © 2018, Juniper Networks, Inc.

Table 34: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
shaping	ATM interfaces	*Defining the ATM Traffic-Shaping Profile Overview*
shaping	Circuit Emulation PICs	*Configuring ATM QoS or Shaping*
sonet-options	SONET/SDH interfaces	*SONET/SDH Interfaces Overview*
source-address-filter *mac-address*	Aggregated Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, and Gigabit Ethernet interfaces	*Configuring MAC Address Filtering for Ethernet Interfaces*
(source-filtering | no-source-filtering)	Aggregated Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, Gigabit Ethernet IQ and IQE, and Gigabit Ethernet interfaces with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router)	*Configuring MAC Address Filtering for Ethernet Interfaces*
speed (10m | 100m | 1g | auto)	Management Ethernet interface (**fxp0** or **em0**), Tri-Rate Ethernet copper interfaces, and 12-port and 48-port Fast Ethernet interfaces	"Configuring the Interface Speed on Ethernet Interfaces" on page 116, "Understanding Management Ethernet Interfaces" on page 10
speed (oc3 | oc12 | oc48)	SONET/SDH PICs with SFP	"Configuring SONET/SDH Interface Speed" on page 120
stacked-vlan-tagging	Gigabit Ethernet IQ interfaces	*Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview*
start-end-flag (filler | shared)	DS0, E1, E3, T1, and T3 interfaces	*Configuring E1 Start and End Flags*, *Configuring the E3 Start and End Flags*, *Configuring T1 Start and End Flags*, and *Configuring T3 Start and End Flags*
switching-mode (bidirectional | unidirectional)	Unchannelized OC3, OC12, and OC48 SONET/SDH interfaces on T Series routers	*Configuring Switching Mode*
syslog	Adaptive services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
(syslog | no-syslog)	Adaptive services, monitoring services, and collector interfaces	"Configuring Multiservice Physical Interface Properties" on page 170
t1-options	T1 interfaces	*T1 Interfaces Overview*
t3-options	T3 interfaces	*T3 Interfaces Overview*

Copyright © 2018, Juniper Networks, Inc.

Table 34: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
t391 *seconds*	Link services and voice services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
t392 *number*	Link services and voice services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
threshold *bytes*	All interfaces, except ATM, channelized E1, E1, Fast Ethernet, Gigabit Ethernet, and channelized IQ	"Configuring Receive and Transmit Leaky Bucket Properties to Reduce Network Congestion" on page 156
timeslots *time-slot-range*	Channelized T1 IQ and channelized E1 IQ interfaces	*Channelized Interfaces Overview*
tm (ignore | normal | require)	Serial interfaces (EIA-530)	"Configuring the Serial Signal Handling" on page 344
tm-polarity (negative | positive)	Serial interfaces (EIA-530)	"Configuring Serial Signal Polarities" on page 347
traceoptions	All	"Tracing Operations of an Individual Router Interface" on page 355
traceoptions	All	"Tracing Operations of the Interface Process" on page 356
transmit-bucket	All interfaces, except ATM, Fast Ethernet, Tri-Rate Ethernet copper, and Gigabit Ethernet	"Configuring Receive and Transmit Leaky Bucket Properties to Reduce Network Congestion" on page 156
transmit-clock invert	Serial interfaces (EIA-530, V.35, and X.21)	"Configuring the Serial Clocking Mode" on page 342
transmit-weight (cells *number* | percent *number*)	ATM2 IQ interfaces	*ATM2 IQ VC Tunnel CoS Components Overview*
(traps | no-traps)	All	"Enabling or Disabling SNMP Notifications on Physical Interfaces" on page 171
trigger *defect* ignore | *defect* hold-time up *milliseconds* down *milliseconds*;	10-Gigabit Ethernet interfaces in WAN PHY mode and ATM over SONET/SDH and SONET/SDH interfaces	*Configuring SONET/SDH Defect Triggers*
(unframed | no-unframed)	E3 IQ interfaces	*Configuring E3 IQ and IQE Unframed Mode*

 Copyright © 2018, Juniper Networks, Inc.

Table 34: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
unidirectional	10-Gigabit Ethernet interfaces on: • MX960 4–Port 10–Gigabit Ethernet DPC • T Series 10–Gigabit Ethernet IQ2 PIC • T Series 10–Gigabit Ethernet IQ2E PIC	"Enabling Unidirectional Traffic Flow on Physical Interfaces" on page 158
vbr peak *rate* sustained *rate* burst *length*	ATM interfaces	*Defining the ATM Traffic-Shaping Profile Overview*
vc-cos-mode (alternate | strict)	ATM2 interfaces	*ATM2 IQ VC Tunnel CoS Components Overview*
vlan-tagging	Fast Ethernet, Tri-Rate Ethernet copper, and Gigabit Ethernet interfaces	*802.1Q VLANs Overview*
vlan-vci-tagging	Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, 10-Gigabit Ethernet, and aggregated Ethernet interfaces	"Configuring ATM-to-Ethernet Interworking" on page 282
vpi *vpi-identifier*	ATM interfaces	*Configuring ATM Cell-Relay Promiscuous Mode* and *Configuring the Maximum Number of ATM1 VCs on a VP*
vtmapping	Channelized STM1 interfaces	*Configuring Virtual Tributary Mapping of Channelized STM1 Interfaces*
warning*warning-name (syslog | link-down)*	10-Gigabit Ethernet interfaces	*Configuring 10-Gigabit Ethernet Link Down Notification for Optics Options Alarm or Warning*
wavelength *nm*	Gigabit Ethernet dense wavelength-division multiplexing (DWDM) interfaces	*Ethernet DWDM Interface Wavelength Overview*
working-circuit *group-name*	SONET/SDH interfaces	*Configuring Switching Between the Working and Protect Circuits*
yellow-differential-delay *milliseconds*	Link services and voice services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
(z0-increment | no-z0-increment)	SONET/SDH interfaces	*Configuring an Incrementing STM ID to Interoperate with Older Equipment in SDH Mode*

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

Configuring Interface Ranges

NOTE: This task uses Junos OS for EX Series switches that does not support the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that supports ELS, see *Configuring Interface Ranges for EX Series Switches with ELS*. For ELS details, see *Using the Enhanced Layer 2 Software CLI*.

The Junos OS allows you to group a range of identical interfaces into an *interface range*. You first specify the group of identical interfaces in the interface range. Then you can apply a common configuration to the specified interface range, reducing the number of configuration statements required and saving time while producing a compact configuration.

Configuring Interface Ranges

To configure an interface range, include the **interface-range** statement at the **[edit interfaces]** hierarchy level.

The **interface-range** statement accepts only physical networking interface names in its definition. The following interface types are supported and example CLI descriptors are shown:

- ATM—**at-*fpc/pic/port***
- Channelized—**(coc | cstm)*n-fpc/pic/port***
- DPC—**xe-*fpc/pic/port***
- E1/E3—**(e1 | e3)-*fpc/pic/port***
- Ethernet—**(xe | ge | fe)-*fpc/pic/port***
- ISDN—**isdn-*fpc/pic/port***
- Serial—**se-*fpc/pic/port***
- SONET/SDH—**so-*fpc/pic/port***
- T1/T3—**(t1 | t3)-*fpc/pic/port***

Copyright © 2018, Juniper Networks, Inc.

Interfaces can be grouped either as a range of interfaces or using a number range under the **interface-range** statement definition.

Interfaces in an **interface-range** definition can be added as part of a member range or as individual members or multiple members using a number range.

To specify a member range, use the **member-range** statement at the **[edit interfaces interface-range *name*]** hierarchy level.

To specify interfaces in lexical order, use the **member-range *start-range* to *end-range*** statement.

A range for a member statement should contain the following:

- *****—All, specifies sequential interfaces from 0 through 47.

> CAUTION: **The wildcard * in a member statement does not take into account the interface numbers supported by a specific interface type. Irrespective of the interface type, * includes interface numbers ranging from 0 through 47 to the interface group. Therefore, use * in a member statement with caution.**

- **num**—Number, specifies one specific interface by its number.
- **[low-high]**—Numbers between low to high, specifies a range of sequential interfaces.
- **[num1, num2, num3]**—Numbers **num1**, **num2**, and **num3** specify multiple specific interfaces.

Example: Specifying an Interface Range Member Range

```
member-range ge-0/0/0 to ge-4/0/40;
```

To specify one or multiple members, use the **member** statement at the **[edit interfaces interface-range *name*]** hierarchy level.

To specify the list of interface range members individually or for multiple interfaces using regex, use the **member *list of interface names*** statement.

Example: Specifying an Interface Range Member

```
member ge-0/0/0;
member ge-0/*/*
member ge-0/[1-10]/0;
member ge-0/[1,2,3]/3;
```

Regex or wildcards are not supported for interface-type prefixes. For example, prefixes **ge**, **fe**, and **xe** must be mentioned explicitly.

An **interface-range** definition can contain both **member** and **member-range** statements within it. There is no maximum limit on the number of **member** or **member-range** statements within an interface-range. However, at least one **member** or **member-range** statement must exist within an **interface-range** definition.

Copyright © 2018, Juniper Networks, Inc.

Example: Interface Range Common Configuration

Configuration common to an interface range can be added as a part of the **interface-range** definition, as follows:

```
[edit]
interfaces {
  + interface-range foo {
  + member-range ge-1/0/0 to ge-4/0/40;
  + member ge-0/1/1;
  + member ge-5/[1-10]/*;
    /*Common configuration is added as part of interface-range definition*/
    mtu 256;
    hold-time up 10;
    ether-options {
      flow-control;
      speed {
        100m;
      }
      802.3ad primary;
    }
  }
}
```

An **interface-range** definition having just **member** or **member-range** statements and no common configurations statements is valid.

These defined interface ranges can be used in other configuration hierarchies, in places where an **interface** node exists.

Example: Interface-Range foo Used Under the Protocols Hierarchy

```
protocols {
  dot1x {
    authenticator {
      interface foo{
        retries 1;
      }
    }
  }
}
```

foo should be an **interface-range** defined at the **[interfaces]** hierarchy level. In the above example, the **interface** node can accept both individual interfaces and interface ranges.

TIP: To view an interface range in expanded configuration, use the **(show | display inheritance)** command. For more information, see the *CLI User Guide*.

By default, **interface-range** is not available to configure in the CLI where the **interface** statement is available. The following locations are supported; however, some of the hierarchies shown in this list are product specific:

- **protocols dot1x authentication interface**
- **protocols dvmrp interface**
- **protocols oam ethernet lmi interface**

 Copyright © 2018, Juniper Networks, Inc.

- protocols esis interface
- protocols igmp interface
- protocols igmp-host client *num* interface
- protocols mld-host client *num* interface
- protocols router-advertisement interface
- protocols isis interface
- protocols ldp interface
- protocols oam ethernet link-fault-management interface
- protocols lldp interface
- protocols link-management peer lmp-control-channel interface
- protocols link-management peer control-channel
- protocols link-management te-link *name* interface
- protocols mld interface
- protocols ospf area *id* interface
- protocols pim interface
- protocols router-discovery interface
- protocols rip group *name* neighbour
- protocols ripng group *name* neighbour
- protocols rsvp interface
- protocols snmp interface
- protocols layer2-control bpdu-block interface
- protocols layer2-control mac-rewrite interface
- protocols mpls interface
- protocols stp interface
- protocols rstp interface
- protocols mstp interface
- protocols vstp interface
- protocols mstp msti *id* interface
- protocols mstp msti vlan *id* interface
- protocols vstp vlan *name* interface
- protocols gvrp interface
- protocols igmp-snooping vlan *name* interface
- protocols lldp interface

- **protocols lldp-med interface**
- **protocols sflow interfaces**
- **ethernet-switching-options analyzer *name* input [egress | ingress] interface**
- **ethernet-switching-options analyzer *name* output interface**
- **ethernet-switching-options secure-access-port interface**
- **ethernet-switching-options interfaces ethernet-switching-options voip interface**
- **ethernet-switching-options redundant-trunk-group group g1 interface**
- **ethernet-switching-options redundant-trunk-group group g1 interface**
- **ethernet-switching-options bpdu-block interface**
- **poe interface vlans pro-bng-mc1-bsd1 interface**

See Also

- Expanding Interface Range Member and Member Range Statements on page 94
- Configuration Inheritance for Member Interfaces on page 95
- Member Interfaces Inheriting Configuration from Configuration Groups on page 97
- Interfaces Inheriting Common Configuration on page 98
- Configuring Inheritance Range Priorities on page 98
- Configuration Expansion Where Interface Range Is Used on page 99
- *Physical Interfaces*

Expanding Interface Range Member and Member Range Statements

All **member** and **member-range** statements in an interface range definition are expanded to generate the final list of interface names for the specified interface range.

Example: Expanding Interface Range Member and Member Range Statements

```
[edit]
  interfaces {
  interface-range range-1 {
    member-range ge-0/0/0 to ge-4/0/20;
    member ge-10/1/1;
    member ge-5/[0-5]/*;
    /*Common configuration is added part of the interface-range definition*/
    mtu 256;
    hold-time up 10;
    ether-options {
      flow-control;
      speed {
        100m;
      }
      802.3ad primary;
    }
  }
```


Copyright © 2018, Juniper Networks, Inc.

```
}
```

For the **member-range** statement, all possible interfaces between **start-range** and **end-range** are considered in expanding the members. For example, the following **member-range** statement:

member-range ge-0/0/0 to ge-4/0/20
expands to:

```
[ge-0/0/0, ge-0/0/1 ... ge-0/0/max_ports
 ge-0/1/0  ge-0/1/1 ... ge-0/1/max_ports
 ge-0/2/0  ge-0/2/1 ... ge-0/2/max_ports
                       .
                       .
 ge-0/MAX_PICS/0 ... ge-0/max_pics/max_ports
 ge-1/0/0  ge-1/0/1 ... ge-1/0/max_ports
                       .
 ge-1/MAX_PICS/0 ... ge-1/max_pics/max_ports
                       .
                       .
 ge-4/0/0 ge-4/0/1  ... ge-4/0/max_ports]
```

The following **member** statement:

ge-5/[0-5]/*
expands to:

```
ge-5/0/0 ... ge-5/0/max_ports
ge-5/1/0 ... ge-5/0/max_ports
    .
     .
ge-5/5/0 ... ge-5/5/max_ports
```

The following **member** statement:

ge-5/1/[2,3,6,10]
expands to:

```
ge-5/1/2
ge-5/1/3
ge-5/1/6
ge-5/1/10
```

See Also
- *Physical Interfaces*

Configuration Inheritance for Member Interfaces

When the Junos OS expands the **member** and **member-range** statements present in an **interface-range**, it creates *interface objects* if they are not explicitly defined in the configuration. The common configuration is copied to all its member interfaces in the **interface-range**.

Copyright © 2018, Juniper Networks, Inc.

Example: Configuration Priorities

Foreground interface configuration takes priority compared to configuration inherited by the interface through the **interface-range**.

```
interfaces {
  interface-range range-1 {
    member-range ge-1/0/0/ to ge-10/0/47;
    mtu 256;
    }
    ge-1/0/1 {
      mtu 1024;
    }
}
```

In the preceding example, interface **ge-1/0/1** will have an MTU value of 1024.

This can be verified with output of the **show interfaces | display inheritance** command, as follows:

```
user@host: # show interfaces | display inheritance
## 'ge-1/0/0' was expanded from interface-range 'range-1'
##
ge-1/0/0 {
    ##
    ## '256' was expanded from interface-range 'range-1'
    ##
    mtu 256;
}
ge-1/0/1 {
    mtu 1024;
}
##
## 'ge-1/0/2' was expanded from interface-range 'range-1'
##
ge-1/0/2 {
    ##
    ## '256' was expanded from interface-range 'range-1'
    ##
    mtu 256;
}
      .........
      .........
##
## 'ge-10/0/47' was expanded from interface-range 'range-1'
##
ge-10/0/47 {
    ##
    ## '256' was expanded from interface-range 'range-1'
    ##
    mtu 256;
}
```

See Also

- *Physical Interfaces*

 Copyright © 2018, Juniper Networks, Inc.

Member Interfaces Inheriting Configuration from Configuration Groups

Interface range member interfaces inherit the config-groups configuration like any other foreground configuration. **interface-range** is similar to any other foreground configuration statement. The only difference is that the **interface-range** goes through a member interfaces expansion before Junos OS reads this configuration.

```
groups {
  global {
    interfaces {
      <*> {
        hold-time up 10;
      }
    }
  }
  apply-groups [global];
  interfaces {
    interface-range range-1 {
      member-range ge-1/0/0 to ge-10/0/47;
      mtu 256;
    }
  }
}
```

The **hold-time** configuration is applied to all members of **interface-range** ***range-1***.

This can be verified with **show interfaces | display inheritance** as follows:

```
user@host# show interfaces | display inheritance
ge-1/0/0 {
    ##
    ## '256' was expanded from interface-range 'range-1'
    ##
    mtu 256;
    ##
    ## 'hold-time' was inherited from group 'global'
    ## '10' was inherited from group 'global'
    ##
    hold-time up 10;
}
ge-1/0/1 {
    ##
    ## '256' was expanded from interface-range 'range-1'
    ##
    mtu 256;
    ##
    ## 'hold-time' was inherited from group 'global'
    ## '10' was inherited from group 'global'
    ##
    hold-time up 10;
}
ge-10/0/47 {
    ##
    ## '256' was expanded from interface-range 'range-1'
    ##
    mtu 256;
    ##
    ## 'hold-time' was inherited from group 'global'
```

Copyright © 2018, Juniper Networks, Inc.

```
    ## '10' was inherited from group 'global'
    ##
    hold-time up 10;
}
```

See Also
- *Using Wildcards with Configuration Groups*

Interfaces Inheriting Common Configuration

If an interface is a member of several interface ranges, that interface will inherit the common configuration from all of those interface ranges.

```
[edit]
interfaces {
  interface-range range-1 {
    member-range ge-1/0/0 to ge-10/0/47;
    mtu 256;
  }
  }
  interfaces {
    interface-range range-1 {
      member-range ge-10/0/0 to ge-10/0/47;
      hold-time up 10;
    }
  }
```

In this example, interfaces **ge-10/0/0** through **ge-10/0/47** will have both **hold-time** and **mtu**.

Configuring Inheritance Range Priorities

The interface ranges are defined in the order of inheritance priority, with the first interface range configuration data taking priority over subsequent interface ranges.

```
[edit]
interfaces {
  interface-range int-grp-one {
    member-range ge-0/0/0 to ge-4/0/40;
    member ge-1/1/1;
    /*Common config is added part of the interface-range definition*/
    mtu 256;
    hold-time up 10;
  }
  }
  interfaces {
    interface-range int-grp-two {
      member-range ge-5/0/0 to ge-10/0/40;
      member ge-1/1/1;
      mtu 1024;
    }
  }
```

Interface **ge-1/1/1** exists in both **interface-range** ***int-grp-one*** and **interface-range** ***int-grp-two***. This interface inherits **mtu** ***256*** from **interface-range** ***int-grp-one*** because it was defined first.

 Copyright © 2018, Juniper Networks, Inc.

See Also
- *Physical Interfaces*

Configuration Expansion Where Interface Range Is Used

In this example, **interface-range** ***range-1*** is used under the **protocols** hierarchy:

```
[edit]
interfaces {
   interface-range range-1 {
      member ge-10/1/1;
      member ge-5/5/1;
      mtu 256;
      hold-time up 10;
      ether-options {
         flow-control;
         speed {
            100m;
         }
         802.3ad primary;
      }
   }
   protocols {
      dot1x {
         authenticator {
            interface range-1 {
               retries 1;
            }
         }
      }
   }
}
```

The **interface** node present under **authenticator** is expanded into member interfaces of the **interface-range** ***range-1*** as follows:

```
protocols {
   dot1x {
      authenticator {
         interface ge-10/1/1 {
            retries 1;
         }
         interface ge-5/5/1 {
            retries 1;
         }
      }
   }
}
```

The **interface** ***range-1*** statement is expanded into two interfaces, ge-10/1/1 and ge-5/5/1, and configuration **retries** ***1*** is copied under those two interfaces.

This configuration can be verified using the **show protocols dot1x | display inheritance** command.

See Also
- *Physical Interfaces*

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- *Physical Interfaces*

Specifying an Aggregated Interface

The M Series, MX Series, and T Series routers support aggregated interfaces. To specify an aggregated interface assign a number with the aggregated interface name. For example, configure **ae*x*** at the **[edit interfaces]** hierarchy level, where *x* is an integer ranging 0 through 127 for M Series and T Series routers and 0 through 479 on MX Series routers.

For aggregated SONET/SDH interfaces, configure **as*x*** at the **[edit interfaces]** hierarchy level.

NOTE: **SONET/SDH aggregation is proprietary to the Junos OS and might not work with other software.**

If you are configuring VLANs for aggregated Ethernet interfaces, you must include the **vlan-tagging** statement at the **[edit interfaces ae*x*]** hierarchy level to complete the association.

Related Documentation

- *Aggregated Ethernet Interfaces Overview*
- *Configuring Aggregated SONET/SDH Interfaces*

Media MTU Overview

The media maximum transmission unit (MTU) is the largest data unit that can be forwarded without fragmentation.

The default media MTU size used on a physical interface depends on the encapsulation used on that interface. In some cases, the default IP Protocol MTU depends on whether the protocol used is IP version 4 (IPv4) or International Organization for Standardization (ISO).

The default media MTU is calculated as follows:

Default media MTU = Default IP MTU + encapsulation overhead

When you are configuring point-to-point connections, the MTU sizes on both sides of the connections must be the same. Also, when you are configuring point-to-multipoint connections, all interfaces in the subnet must use the same MTU size.

 Copyright © 2018, Juniper Networks, Inc.

NOTE: The actual frames transmitted also contain cyclic redundancy check (CRC) bits, which are not part of the media MTU. For example, the media MTU for a Gigabit Ethernet Version 2 interface is specified as 1514 bytes, but the largest possible frame size is actually 1518 bytes; you need to consider the extra bits in calculations of MTUs for interoperability.

The physical MTU for Ethernet interfaces does not include the 4-byte frame check sequence (FCS) field of the Ethernet frame.

A SONET/SDH interface operating in concatenated mode has a “c” added to the rate descriptor. For example, a concatenated OC48 interface is referred to as OC48c.

If you do not configure an MPLS MTU, the Junos OS derives the MPLS MTU from the physical interface MTU. From this value, the software subtracts the encapsulation-specific overhead and space for the maximum number of labels that might be pushed in the Packet Forwarding Engine. Currently, the software provides for three labels of four bytes each, for a total of 12 bytes.

In other words, the formula used to determine the MPLS MTU is the following:

MPLS MTU = physical interface MTU – encapsulation overhead – 12

Related Documentation

- Configuring the Media MTU on page 115
- Encapsulation Overhead by Interface Encapsulation Type on page 112
- Encapsulation Overhead by Interface Encapsulation Type on page 112

Media MTU Sizes by Interface Type

The media maximum transmission unit (MTU) is the largest data unit that can be forwarded without fragmentation.

If you change the size of the media MTU, you must ensure that the size is equal to or greater than the sum of the protocol MTU and the encapsulation overhead.

This topic includes following information:

Copyright © 2018, Juniper Networks, Inc.

- Media MTU Sizes by Interface Type for EX Series Switches and ACX Series Routers on page 108
- Media MTU Sizes by Interface Type for PTX Series Packet Transport Routers on page 108

Media MTU Sizes by Interface Type for M5 and M7i Routers with CFEB, M10 and M10i Routers with CFEB, and M20 and M40 Routers

Table 35: Media MTU Sizes by Interface Type for M5 and M7i Routers with CFEB, M10 and M10i Routers with CFEB, and M20 and M40 Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Adaptive Services (MTU size not configurable)	9192	N/A	N/A
ATM	4482	9192	4470
E1/T1	1504	9192	1500
E3/T3	4474	9192	4470
Fast Ethernet	1514	1533 (4-port) 1532 (8-port) 1532 (12-port) NOTE: The maximum MTU for two 100Base-TX Fast Ethernet port FIC is 9192 bytes.	1500 (IPv4), 1497 (ISO)
Gigabit Ethernet	1514	9192 NOTE: The maximum MTU for one Gigabit Ethernet port FIC is 9192 bytes.	1500 (IPv4), 1497 (ISO)
Serial	1504	9192	1500 (IPv4), 1497 (ISO)
SONET/SDH	4474	9192	4470

 Copyright © 2018, Juniper Networks, Inc.

Media MTU Sizes by Interface Type for M40e Routers

Table 36: Media MTU Sizes by Interface Type for M40e Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Adaptive Services (MTU size not configurable)	9192	N/A	N/A
ATM	4482	9192	4470
E1/T1	1504	4500	1500
E3/T3	4474	4500 9192 (4-port)	4470
E3/DS3 IQ	4474	9192	4470
Fast Ethernet	1514	1533	1500 (IPv4), 1497 (ISO)
Gigabit Ethernet	1514	9192 (1- or 2-port) 9192 (4-port)	1500 (IPv4), 1497 (ISO)
Serial	1504	9192	1500 (IPv4), 1497 (ISO)
SONET/SDH	4474	4500 (1-port nonconcatenated) 9192 (4-port OC3) 9192 (4-port OC3c) 4500 (1-port OC12) 4500 (4-port OC12) 4500 (4-port OC12c) 4500 (1-port OC48) 9192 (2-port OC3) 9192 (2-port OC3c) 9192 (1-port OC12c) 9192 (1-port OC48c) 4500 (1-port OC192) 9192 (1-port OC192c)	4470

Copyright © 2018, Juniper Networks, Inc.

Media MTU Sizes by Interface Type for M160 Routers

Table 37: Media MTU Sizes by Interface Type for M160 Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Adaptive Services (MTU size not configurable)	9192	N/A	N/A
ATM	4482	9192	4470
E1/T1	1504	4500	1500
E3/T3	4474	4500	4470
E3/DS3 IQ	4474	9192	4470
Fast Ethernet	1514	1533	1500 (IPv4), 1497 (ISO)
Gigabit Ethernet	1514	9192 (1- or 2-port) 4500 (4-port)	1500 (IPv4), 1497 (ISO)
Serial	1504	9192	1500 (IPv4), 1497 (ISO)
SONET/SDH	4474	4500 (1-port nonconcatenated) 9192 (1- or 2-port) 4500 (4-port)	4470

Media MTU Sizes by Interface Type for M7i Routers with CFEB-E, M10i Routers with CFEB-E, and M320 and M120 Routers

Table 38: Media MTU Sizes by Interface Type for M7i Routers with CFEB-E, M10i Routers with CFEB-E, and M320 and M120 Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
ATM2 IQ	4482	9192	4470
Channelized DS3 IQ	4471	4500	4470
Channelized E1 IQ	1504	4500	1500
Channelized OC12 IQ	4474	9192	4470

Copyright © 2018, Juniper Networks, Inc.

Table 38: Media MTU Sizes by Interface Type for M7i Routers with CFEB-E, M10i Routers with CFEB-E, and M320 and M120 Routers (continued)

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Channelized STM1 IQ	4474	9192	4470
DS3	4471	4500	4470
E1	1504	4500	1500
E3 IQ	4471	4500	4470
Fast Ethernet	1514	1533 (4-port) 1532 (8-, 12- and 48-port)	1500 (IPv4), 1497 (ISO)
Gigabit Ethernet	1514	9192	1500 (IPv4), 1497 (ISO)
SONET/SDH	4474	9192	4470
T1	1504	4500	1500
CT3 IQ (excluding M120)	4474	9192	4470

Media MTU Sizes by Interface Type for MX Series Routers

Table 39: Media MTU Sizes by Interface Type for MX Series Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Gigabit Ethernet	1514	• 9192 • 9500 (Junos OS 16.1R1 and later releases)	1500 (IPv4), 1488 (MPLS), 1497 (ISO)
10-Gigabit Ethernet	1514	• 9192 • 9500 (Junos OS 16.1R1 and later releases)	1500 (IPv4), 1488 (MPLS), 1497 (ISO)
Multi-Rate Ethernet	1514	• 9192 • 9500 (Junos OS 16.1R1 and later releases)	1500 (IPv4), 1488 (MPLS), 1497 (ISO)
Tri-Rate Ethernet	1514	• 9192 • 9500 (Junos OS 16.1R1 and later releases)	1500 (IPv4), 1488 (MPLS), 1497 (ISO)

Copyright © 2018, Juniper Networks, Inc.

Table 39: Media MTU Sizes by Interface Type for MX Series Routers (continued)

Channelized SONET/SDH OC3/STM1 (Multi-Rate)	1514	9192	1500 (IPv4), 1488 (MPLS), 1497 (ISO)
DS3/E3 (Multi-Rate)	1514	9192	1500 (IPv4), 1488 (MPLS), 1497 (ISO)

NOTE: Starting in Junos OS Release 16.1R1, the MTU size for a media or protocol is increased from 9192 to 9500 for Ethernet interfaces on the following MX Series MPCs:

- MPC1
- MPC2
- MPC2E
- MPC3E
- MPC4E
- MPC5E
- MPC6E

NOTE: Starting in Junos OS Release 16.1R1, the MTU size for a media or protocol is increased from 9192 to 9500 for Ethernet interfaces on the following MX Series MPCs:

- MPC1
- MPC2
- MPC2E
- MPC3E
- MPC4E
- MPC5E
- MPC6E

Starting in Junos OS Release 16.1R1, the MTU size has been increased to 16,000 bytes for certain MPCs. The MTU size for the following MPCs has been increased to 16000 bytes:

- MPC7E (MPC7E-MRATE and MP7E-10G)
- MPC8E (MX2K-MPC8E)
- MPC9E (MX2K-MPC9E)

Starting in Junos OS Release 17.3R1, the MTU size for MX10003 MPC is 16,000 bytes.

Starting in Junos OS Release 17.4R1, the MTU size for MX204 is 16,000 bytes.

 Copyright © 2018, Juniper Networks, Inc.

Media MTU Sizes by Interface Type for T320 Routers

Table 40: Media MTU Sizes by Interface Type for T320 Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
ATM	4482	9192	4470
ATM2 IQ	4482	9192	4470
Channelized OC12 IQ	4474	9192	4470
Channelized STM1 IQ	4474	9192	4470
DS3	4471	4500	4470
Fast Ethernet	1514	1533 (4-port) 1532 (12- and 48-port)	1500 (IPv4), 1497 (ISO)
Gigabit Ethernet	1514	9192	1500 (IPv4), 1497 (ISO)
SONET/SDH	4474	9192	4470
CT3 IQ	4474	9192	4470

Media MTU Sizes by Interface Type for T640 Platforms

Table 41: Media MTU Sizes by Interface Type for T640 Platforms

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
ATM2 IQ	4482	9192	4470
48-port Fast Ethernet	1514	1532	1500 (IPv4), 1497 (ISO)
Gigabit Ethernet	1514	9192	1500 (IPv4), 1497 (ISO)
SONET/SDH	4474	9192	4470
CT3 IQ	4474	9192	4470

Copyright © 2018, Juniper Networks, Inc.

Media MTU Sizes by Interface Type for EX Series Switches and ACX Series Routers

Table 42: Media MTU Sizes by Interface Type for EX Series Switches and ACX Series Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Gigabit Ethernet	1514	9192	1500 (IPv4), 1497 (ISO)
10-Gigabit Ethernet	1514	9192	1500 (IPv4), 1497 (ISO)

NOTE: On ACX Series routers, you can configure the protocol MTU by including the **mtu** statement at the **[edit interfaces** *interface-name* **unit** *logical-unit-number* **family inet]** or **[edit interfaces** *interface-name* **unit** *logical-unit-number* **family inet6]** hierarchy level.

- If you configure the protocol MTU at any of these hierarchy levels, the configured value is applied to all families that are configured on the logical interface.
- If you are configuring the protocol MTU for both **inet** and **inet6** families on the same logical interface, you must configure the same value for both the families. It is not recommended to configure different MTU size values for **inet** and **inet6** families that are configured on the same logical interface.

Media MTU Sizes by Interface Type for PTX Series Packet Transport Routers

Table 43: Media MTU Sizes by Interface Type for PTX Series Packet Transport Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
10-Gigabit Ethernet	1514	9500	1500 (IPv4), 1488 (MPLS), 1497 (ISO)
40-Gigabit Ethernet	1514	9500	1500 (IPv4), 1488 (MPLS), 1497 (ISO)
100-Gigabit Ethernet	1514	9500	1500 (IPv4), 1488 (MPLS), 1497 (ISO)

 Copyright © 2018, Juniper Networks, Inc.

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, the MTU size for MX204 is 16,000 bytes.
17.3R1	Starting in Junos OS Release 17.3R1, the MTU size for MX10003 MPC is 16,000 bytes.
16.1R1	Starting in Junos OS Release 16.1R1, the MTU size has been increased to 16,000 bytes for certain MPCs.

Related Documentation

- Encapsulation Overhead by Interface Encapsulation Type on page 112
- Configuring the Media MTU on page 115
- Media MTU Overview on page 100
- Setting the Protocol MTU on page 224

Configuring the Media MTU on ACX Series Routers

Media MTU Overview

The default media MTU size used on a physical interface depends on the encapsulation used on that interface. In some cases, the default IP Protocol MTU depends on whether the protocol used is IP version 4 (IPv4) or International Organization for Standardization (ISO).

The default media MTU is calculated as follows:

Default media MTU = Default IP MTU + encapsulation overhead

When you are configuring point-to-point connections, the MTU sizes on both sides of the connections must be the same. Also, when you are configuring point-to-multipoint connections, all interfaces in the subnet must use the same MTU size. For details about encapsulation overhead, see “Encapsulation Overhead by Encapsulation Type” on page 111.

NOTE: The actual frames transmitted also contain cyclic redundancy check (CRC) bits, which are not part of the media MTU. For example, the media MTU for a Gigabit Ethernet Version 2 interface is specified as 1514 bytes, but the largest possible frame size is actually 1518 bytes; you need to consider the extra bits in calculations of MTUs for interoperability.

The physical MTU for Ethernet interfaces does not include the 4-byte frame check sequence (FCS) field of the Ethernet frame.

If you do not configure an MPLS MTU, the Junos OS derives the MPLS MTU from the physical interface MTU. From this value, the software subtracts the encapsulation-specific overhead and space for the maximum number of labels that might be pushed in the Packet Forwarding Engine. Currently, the software provides for three labels of four bytes each, for a total of 12 bytes.

In other words, the formula used to determine the MPLS MTU is the following:

MPLS MTU = physical interface MTU – encapsulation overhead – 12

If you configure an MTU value by including the **mtu** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family mpls]** hierarchy level, the configured value is used.

How to Configure the Media MTU

To modify the default media MTU size for a physical interface, include the **mtu** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
mtu bytes;
```

If you change the size of the media MTU, you must ensure that the size is equal to or greater than the sum of the protocol MTU and the encapsulation overhead.

NOTE: Changing the media MTU or protocol MTU causes an interface to be deleted and added again.

You configure the protocol MTU by including the **mtu** statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* family inet]**
- **[edit interfaces *interface-name* unit *logical-unit-number* family inet6]**

If you configure the protocol MTU at any of these hierarchy levels, the configured value is applied to all families that are configured on the logical interface.

 Copyright © 2018, Juniper Networks, Inc.

NOTE: If you are configuring the protocol MTU for both **inet** and **inet6** families on the same logical interface, you must configure the same value for both the families. It is not recommended to configure different MTU size values for **inet** and **inet6** families that are configured on the same logical interface.

The ACX Series routers do not support per-family maximum transmission unit (MTU) configuration. The MTU applied to family **inet** gets applied to other families as well, even though it can be configured though CLI and visible in the **show interface extensive** output.

Encapsulation Overhead by Encapsulation Type

Table 44: Encapsulation Overhead by Encapsulation Type

Interface Encapsulation	Encapsulation Overhead (Bytes)
802.1Q/Ethernet 802.3	21
802.1Q/Ethernet Subnetwork Access Protocol (SNAP)	26
802.1Q/Ethernet version 2	18
ATM Cell Relay	4
ATM permanent virtual connection (PVC)	12
Cisco HDLC	4
Ethernet 802.3	17
Ethernet circuit cross-connect (CCC) and virtual private LAN service (VPLS)	4
Ethernet over ATM	32
Ethernet SNAP	22
Ethernet translational cross-connect (TCC)	18
Ethernet version 2	14
Extended virtual local area network (VLAN) CCC and VPLS	4
Extended VLAN TCC	22
Frame Relay	4
PPP	4

Copyright © 2018, Juniper Networks, Inc.

Table 44: Encapsulation Overhead by Encapsulation Type (continued)

Interface Encapsulation	Encapsulation Overhead (Bytes)
VLAN CCC	4
VLAN VPLS	4
VLAN TCC	22

Media MTU Sizes by Interface Type for ACX Series Routers

Table 45: Media MTU Sizes by Interface Type for ACX Series Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Gigabit Ethernet	1514	9192	1500 (IPv4), 1497 (ISO)
10-Gigabit Ethernet	1514	9192	1500 (IPv4), 1497 (ISO)

Related Documentation

- Configuring Interface Encapsulation on Physical Interfaces on page 129
- Setting the Protocol MTU on page 224

Encapsulation Overhead by Interface Encapsulation Type

If you change the size of the media MTU, you must ensure that the size is equal to or greater than the sum of the protocol MTU and the encapsulation overhead. The following table lists the interface encapsulation and corresponding encapsulation overhead.

Table 46: Encapsulation Overhead by Encapsulation Type

Interface Encapsulation	Encapsulation Overhead (Bytes)
802.1Q/Ethernet 802.3	21
802.1Q/Ethernet Subnetwork Access Protocol (SNAP)	26
802.1Q/Ethernet version 2	18
ATM Cell Relay	4
ATM permanent virtual connection (PVC)	12
Cisco HDLC	4
Ethernet 802.3	17

Copyright © 2018, Juniper Networks, Inc.

Table 46: Encapsulation Overhead by Encapsulation Type (continued)

Interface Encapsulation	Encapsulation Overhead (Bytes)
Ethernet circuit cross-connect (CCC) and virtual private LAN service (VPLS)	4
Ethernet over ATM	32
Ethernet SNAP	22
Ethernet translational cross-connect (TCC)	18
Ethernet version 2	14
Extended virtual local area network (VLAN) CCC and VPLS	4
Extended VLAN TCC	22
Frame Relay	4
PPP	4
VLAN CCC	4
VLAN VPLS	4
VLAN TCC	22

Related Documentation

- Media MTU Overview on page 100
- Configuring the Media MTU on page 115
- Encapsulation Overhead by Interface Encapsulation Type on page 112
- Setting the Protocol MTU on page 224

Configuring Interface Description

You can include a text description of each physical interface in the configuration file. Any descriptive text you include is displayed in the output of the **show interfaces** commands, and is also exposed in the **ifAlias** Management Information Base (MIB) object. It has no impact on the interface's configuration.

To add a text description, include the **description** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit]
user@host# set interfaces interface-name description text
```

For example:

```
[edit]
user@host# set interfaces fe-0/0/1 description "Backbone connection to PHL01"
```

The description can be a single line of text. If the text contains spaces, enclose it in quotation marks.

NOTE: You can configure the extended DHCP relay to include the interface description in the option 82 Agent Circuit ID suboption. See *Using DHCP Relay Agent Option 82 Information* in the *Junos OS Broadband Subscriber Management and Services Library*.

For information about describing logical units, see "Adding a Logical Unit Description to the Configuration" on page 188.

To display the description from the router or switch CLI, use the **show interfaces** command:

```
user@host> show interfaces fe-0/0/1
Physical interface: fe-0/0/1, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 23
  Description: Backbone connection to PHL01
  ...
```

To display the interface description from the interfaces MIB, use the **snmpwalk** command from a server. To isolate information for a specific interface, search for the interface index shown in the **SNMP ifIndex** field of the **show interfaces** command output. The **ifAlias** object is in **ifXTable**.

```
user-server> snmpwalk host-fxp0.mylab public ifXTable | grep -e '\.23'
 snmpwalk host-fxp0.mylab public ifXTable | grep -e '\.23'
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifName.23 = fe-0/0/1
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifInMulticastPkts.23 = Counter32: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifInBroadcastPkts.23 = Counter32: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifOutMulticastPkts.23 = Counter32: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifOutBroadcastPkts.23 = Counter32: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifHCInOctets.23 = Counter64: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifHCInUcastPkts.23 = Counter64: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifHCInMulticastPkts.23 = Counter64: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifHCInBroadcastPkts.23 = Counter64: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifHCOutOctets.23 = Counter64: 42
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifHCOutUcastPkts.23 = Counter64: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifHCOutMulticastPkts.23 = Counter64: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifHCOutBroadcastPkts.23 = Counter64: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifLinkUpDownTrapEnable.23 = enabled(1)
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifHighSpeed.23 = Gauge32: 100
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifPromiscuousMode.23 = false(2)
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifConnectorPresent.23 = true(1)
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifAlias.23 = Backbone connection to PHL01
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifCounterDiscontinuityTime.23 = Timeticks:
(0) 0:00:00.00
```

 Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- *Using DHCP Relay Agent Option 82 Information*
- Adding a Logical Unit Description to the Configuration on page 188

Configuring the Media MTU

The media maximum transmission unit (MTU) is the largest data unit that can be forwarded without fragmentation. The default media MTU size used on a physical interface depends on the encapsulation being used on that interface. For a listing of MTU sizes for each encapsulation type, see "Media MTU Sizes by Interface Type" on page 101.

To configure the media-MTU size:

1. In configuration mode, go to the **[edit interfaces *interface-name*]** hierarchy level.

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

2. Include the **mtu** statement.

   [edit interfaces *interface-name*]
   mtu *bytes*;

- If you change the size of the media MTU, you must ensure that the size is equal to or greater than the sum of the protocol MTU and the encapsulation overhead. You configure the protocol MTU by including the **mtu** statement at the following hierarchy levels:
  - **[edit interfaces *interface-name* unit *logical-unit-number* family *family*]**

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

NOTE:

- Changing the media MTU or protocol MTU causes an interface to be deleted and added again.
- Because tunnel services interfaces are considered logical interfaces, you cannot configure the MTU setting for the physical interface. This means you cannot include the mtu statement at the [edit interfaces interface-name] hierarchy level for the following interface types: generic routing encapsulation (gr-), IP-IP (ip-), loopback (lo-), link services (ls-), multilink services (ml-), and multicast (pe-, pd-). You can, however, configure the protocol MTU on all tunnel interfaces except virtual tunnel (vt) interfaces. Starting in Junos OS Release 17.1R3, you cannot configure the maximum transmission unit (MTU) size for vt interfaces because the **mtu bytes** option is deprecated for vt interfaces. Junos OS sets the MTU size for vt interfaces by default to unlimited.
- If you configure an MTU value by including the mtu statement at the [edit interfaces interface-name unit logical-unit-number family mpls] hierarchy level, the configured value is used.

Related Documentation

- Media MTU Overview on page 100
- Media MTU Sizes by Interface Type on page 101
- Encapsulation Overhead by Interface Encapsulation Type on page 112
- Setting the Protocol MTU on page 224

Configuring the Interface Speed

You can configure the interface speed in following ways:

- Configuring the Interface Speed on Ethernet Interfaces on page 116
- Configuring Aggregated Ethernet Link Speed on page 118
- Configuring SONET/SDH Interface Speed on page 120

Configuring the Interface Speed on Ethernet Interfaces

For M Series and T Series Fast Ethernet 12-port and 48-port PIC interfaces, the management Ethernet interface (**fxp0** or **em0**), and the MX Series Tri-Rate Ethernet copper interfaces, you can explicitly set the interface speed. The Fast Ethernet, **fxp0**, and **em0** interfaces can be configured for 10 Mbps or 100 Mbps **(10m | 100m)**. The MX Series Tri-Rate Ethernet copper interfaces can be configured for 10 Mbps, 100 Mbps, or 1 Gbps **(10m | 100m | 1g)**. For information about management Ethernet interfaces and to determine the management Ethernet interface type for your router, see "Understanding

 Copyright © 2018, Juniper Networks, Inc.

Management Ethernet Interfaces" on page 10 and "Supported Routing Engines by Router" on page 17MX Series routers, with MX-DPC and Tri-Rate Copper SFPs, support 20x1 Copper to provide backwards compatibility with 100/10BASE-T and 1000BASE-T operation through an Serial Gigabit Media Independent Interface (SGMII) interface.

1. In configuration mode, go to the **[edit interfaces interface-name]** hierarchy level.

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

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

   [edit interfaces *interface-name*]
   user@host# set speed (10m | 100m | 1g | auto | auto-10m-100m);

NOTE:

- By default, the M Series and T Series routers management Ethernet interface autonegotiates whether to operate at 10 megabits per second (Mbps) or 100 Mbps. All other interfaces automatically choose the correct speed based on the PIC type and whether the PIC is configured to operate in multiplexed mode (using the **no-concatenate** statement in the **[edit chassis]** configuration hierarchy.

- Starting with Junos OS Release 14.2 the **auto-10m-100m** option allows the fixed tri-speed port to auto negotiate with ports limited by **100m** or **10m**maximum speed. This option must be enabled only for Tri-rate MPC port, that is, 3D 40x 1GE (LAN) RJ45 MIC on MX platform. This option does not support other MICs on MX platform.,

- When you manually configure Fast Ethernet interfaces on the M Series and T Series routers, link mode and speed must both be configured. If both these values are not configured, the router uses autonegotiation for the link and ignores the user-configured settings.

- If the link partner does not support autonegotiation, configure either Fast Ethernet port manually to match its link partner's speed and link mode. When the link mode is configured, autonegotiation is disabled.

- On MX Series routers with tri-rate copper SFP interfaces, if the port speed is negotiated to the configured value and the negotiated speed and interface speed do not match, the link will not be brought up.

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

- Starting with Junos OS Release 11.4, half-duplex mode is not supported on Tri-Rate Ethernet copper interfaces. When you include the **speed** statement, you must include the **link-mode full-duplex** statement at the same hierarchy level.

Copyright © 2018, Juniper Networks, Inc.

See Also

- speed on page 979
- *Ethernet Interfaces Overview*
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Aggregated Ethernet Link Speed

On aggregated Ethernet interfaces, you can set the required link speed for all interfaces included in the bundle. Generally, all interfaces that make up a bundle must have the same speed. If you include in the aggregated Ethernet interface an individual link that has a speed different from the speed that you specify in the **link-speed** parameter, an error message is logged. However, there are exceptions.

Starting with Junos OS Release 13.2, aggregated Ethernet supports mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers. For example, these mixes are supported:

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

Starting with Junos OS Release 14.1R1 and 14.2, support for mixed rates on aggregated Ethernet bundles is extended to MX240, MX480, MX960, MX2010, and MX2020 routers.

Starting with Junos OS Release 14.2, aggregated Ethernet supports mixed link speeds on PTX Series Packet Transport Routers.

NOTE:

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

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

Copyright © 2018, Juniper Networks, Inc.

To set the required link speed:

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

   user@host# **edit interfaces** ***interface-name*** aggregated-ether-options

2. Configure the link speed.

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

   speed can be in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

   Aggregated Ethernet interfaces on the M120 router can have one of the following speeds:

   - **100m**—Links are 100 Mbps.
   - **10g**—Links are 10 Gbps.
   - **1g**—Links are 1 Gbps.
   - **oc192**—Links are OC192 or STM64c.

   Aggregated Ethernet links on EX Series switches can be configured to operate at one of the following speeds:

   - **10m**—Links are 10 Mbps.
   - **100m**—Links are 100 Mbps.
   - **1g**—Links are 1 Gbps.
   - **10g**—Links are 10 Gbps.
   - **50g**—Links are 50 Gbps.

   Aggregated Ethernet links on T Series, MX Series, PTX Series routers, and QFX5100, QFX10002, QFX10008, and QFX10016 switches can be configured to operate at one of the following speeds:

   - **100g**—Links are 100 Gbps.
   - **100m**—Links are 100 Mbps.
   - **10g**—Links are 10 Gbps.
   - **1g**—Links are 1 Gbps.
   - **40g**—Links are 40 Gbps.
   - **50g**—Links are 50 Gbps.
   - **80g**—Links are 80 Gbps.
   - **8g**—Links are 8 Gbps.
   - **mixed**—Links are of various speeds.
   - **oc192**—Links are OC192.

Copyright © 2018, Juniper Networks, Inc.

See Also
- aggregated-ether-options on page 432
- *Configuring Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles*
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring SONET/SDH Interface Speed

To configure the speed of SONET/SDH interfaces in concatenated mode:

1. In configuration mode, go to the **[edit interfaces *interface-name*]** hierarchy level, where the *interface-name* is **so-*fpc/pic/port***.

   [edit]
   user@host# **edit interfaces so-*fpc/pic/port***

2. Configure interface speed in concatenated mode.

   For example, each port of 4-port OC12 PIC can be configured to be in OC3 or OC12 speed independently when this PIC is in 4xOC12 concatenated mode.

   [edit interfaces so-*fpc/pic/port*]
   user@host# **set speed (oc3 | oc12 | oc48)**

To configure the speed of SONET/SDH interfaces in nonconcatenated mode:

1. In configuration mode, go to the **[edit interfaces *interface-name*]** hierarchy level, where the ***interface-name*** is **so-*fpc/pic/port***.

   [edit]
   user@host# **edit interfaces so-*fpc/pic/port***

2. Configure interface speed in nonconcatenated mode.

   For example, each port of 4-port OC12 PIC can be configured to be in OC3 or OC12 speed independently when this PIC is in 4xOC12 concatenated mode.

   [edit interfaces so-*fpc/pic/port*]
   user@host# **set speed (oc3 | oc12)**

To configure the PIC to operate in channelized (multiplexed) mode:

1. In configuration mode, go to the **[edit chassis fpc *slot-number* pic *pic-number*]** hierarchy level.

   [edit]
   user@host# **[edit chassis fpc *slot-number* pic *pic-number*]**

2. Configure the **no-concatenate** option.

   [edit interfaces so-*fpc/pic/port*]
   user@host# **set no-concatenate**

 Copyright © 2018, Juniper Networks, Inc.

NOTE: On SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP, Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP, and Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP, you cannot set the interface speed at the **[edit interfaces]** hierarchy level. To enable the speed on these MICs, you need to set the port speed at the **[edit chassis fpc *slot-number* pic *pic-number* port *port-number*]** hierarchy level.

For more information about using the **non-concatenate** statement, see the *Junos OS Administration Library*.

See Also

- *Configuring SONET/SDH Physical Interface Properties*
- *SONET/SDH Interface Speed Overview*
- *SONET/SDH Interfaces Overview*

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2 the **auto-10m-100m** option allows the fixed tri-speed port to auto negotiate with ports limited by **100m** or **10m**maximum speed. This option must be enabled only for Tri-rate MPC port, that is, 3D 40x 1GE (LAN) RJ45 MIC on MX platform. This option does not support other MICs on MX platform.
14.2	Starting with Junos OS Release 14.2, aggregated Ethernet supports mixed link speeds on PTX Series Packet Transport Routers.
14.1	Starting with Junos OS Release 14.1R1 and 14.2, support for mixed rates on aggregated Ethernet bundles is extended to MX240, MX480, MX960, MX2010, and MX2020 routers.
13.2	Starting with Junos OS Release 13.2, aggregated Ethernet supports mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers.
13.2	Starting with Junos OS Release 13.2, 100-Gigabit Ethernet member links can be configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP.
11.4	Starting with Junos OS Release 11.4, half-duplex mode is not supported on Tri-Rate Ethernet copper interfaces. When you include the **speed** statement, you must include the **link-mode full-duplex** statement at the same hierarchy level.

Related Documentation

-

Configuring the Link Characteristics

By default, the router's management Ethernet interface, **fxp0** or **em0**, autonegotiates whether to operate in full-duplex or half-duplex mode. Fast Ethernet interfaces can operate in either full-duplex or half-duplex mode, and all other interfaces can operate only in full-duplex mode. For Gigabit Ethernet, the link partner must also be set to full duplex.

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

NOTE: When you manually configure Fast Ethernet interfaces on the M Series and T Series routers, link mode and speed must both be configured. If both these values are not configured, the router uses autonegotiation for the link and ignores the user-configured settings.

NOTE: When the Fast Ethernet interface on Juniper Networks routers with autonegotiation enabled interoperates with a device configured to operate in half-duplex mode (autonegotiation disabled), the interface defaults to half-duplex mode after the PIC is taken offline and brought back online. This results in packet loss and cyclic redundancy check (CRC) errors.

To explicitly configure an Ethernet interface to operate in either full-duplex or half-duplex mode, include the **link-mode** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
link-mode (full-duplex | half-duplex);
```

Interface Alias Names Overview

You can configure a textual description of a logical unit on a physical interface to be the alias of an interface name. Interface aliasing is supported only at the unit level. If you configure an alias name, the alias name is displayed instead of the interface name in the output of all **show**, **show interfaces**, and other operational mode commands. In Junos OS Release 12.3R8 and later, display of the alias can be suppressed in favor of the actual interface name by using the **display no-interface-alias** parameter along with the show command. Configuring an alias for a logical unit of an interface has no effect on how the interface on the router or switch operates.

When you configure the alias name of an interface, the CLI saves the alias name as the value of the ***interface-name*** variable in the configuration database. To enable backward compatibility with Junos OS releases in which the support for interface aliases is not available, when the Junos OS processes query the configuration database for the ***interface-name*** variable, the actual, exact value of the ***interface-name*** variable is returned instead of the alias name for system operations and computations.

This capability to define interface alias names for physical and logical interfaces is useful in a Junos Node Unifier (JNU) environment that contains a Juniper Networks MX Series

 Copyright © 2018, Juniper Networks, Inc.

3D Universal Edge Router as a controller and EX Series Ethernet switches, QFX Series devices, and ACX Series Universal Metro Routers as satellite devices. The following are the benefits of configuring an alias name, which enables a meaningful, single, and easily identifiable name to be allocated to an interface:

- You can group physical interfaces as one aggregated interface (link aggregation group or LAG bundle) and name that bundle as a satellite connection interface (for example, sat1).
- You can select a logical interface as a member of the LAG bundle or the entire LAG, and name that interface to represent a satellite device port or a service instance (for example, ge-0/0/1).
- You can combine the satellite name and the interface name aliases to wholly represent the satellite port name (for example, sat1:ge-0/0/1 or ge-sat1/0/0/1 or ge-1/0/0/1) in the most easily distinguishable format that denotes a combination of port and satellite parts of the name.

To specify an interface alias, you can use the **alias** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** and **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]** hierarchy levels.

NOTE: In Juniper Networks M Series Multiservice Edge Routers, if the same alias name is configured on more than one logical interface, the router displays an error message and commit fails.

Related Documentation

- Example: Adding an Interface Alias Name on page 123
- **alias on page 435**

Example: Adding an Interface Alias Name

This example shows how to add an alias to the logical unit of an interface. Using an alias to identify interfaces as they appear in the output for operational commands can allow for more meaningful naming conventions and easier identification.

Requirements

This example uses the following hardware and software components:

- One MX Series router that acts as a controller
- One EX4200 switch that acts as a satellite device

- Junos OS Release 13.3R1 or later

Overview

You can create an alias for each logical unit on a physical interface. The descriptive text you define for the alias is displayed in the output of the **show interfaces** commands. In Junos OS Release 12.3R8 and later, display of the alias can be suppressed in favor of the actual interface name by using the **display no-interface-alias** parameter along with the show command. The alias configured for a logical unit of an interface has no effect on how the interface on the router or switch operates – it is only a cosmetic label.

Configuration

Consider a scenario in which alias names are configured on the interfaces of a JNU controller that are connected to a satellite, sat1, in the downlink direction in the JNU management network by using two links. The alias names enable effective, streamlined identification of these interfaces in the operational mode commands that are run on the controller and satellites.

- Configuring Alias Names for the Controller Interfaces on page 124
- Results on page 125

CLI Quick Configuration

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

```
set interfaces ae0 unit 0 alias "controller-sat1-downlink1"
set interfaces ae0.0 family inet address 10.0.0.1/24
set interfaces ae1 unit 0 alias "controller-sat1-downlink1"
set interfaces ae0.0 family inet address 192.0.2.128/25
set interfaces ge-0/0/0 vlan-tagging
set interfaces ge-0/0/0 unit 0 alias "ge-to-corp-gw1"
set interfaces ge-0/0/0.0 vlan-id 101
set interfaces ge-0/0/0.0 family inet address 1.1.1.1/23
set interfaces ge-0/1/0 gigether-options 802.3ad ae0
set interfaces ge-0/1/1 gigether-options 802.3ad ae0
set protocols rip group corporate-firewall neighbor ge-to-corp-gw1
```

Configuring Alias Names for the Controller Interfaces

Step-by-Step Procedure

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

To add an alias name to the controller interfaces that are used to connect to the satellite devices in the downlink direction:

1. Configure an alias name for the logical unit of an aggregated Ethernet interface that is used to connect to a satellite, sat1, in the downlink direction. Configure **inet** family and address for the interface.

 Copyright © 2018, Juniper Networks, Inc.

[edit]
user@host# **set interfaces ae0 unit 0 alias "controller-sat1-downlink1"**
user@host# **set interfaces ae0.0 family inet address 10.0.0.1/24**

2. Configure an alias name for the logical unit of another aggregated Ethernet interface that is used to connect to the same satellite, sat1, in downlink direction. Configure INET family and address for the interface.

   [edit]
   user@host# **set interfaces ae0 unit 1 alias "controller-sat1-downlink2"**
   user@host# **set interfaces ae0.0 family inet address 10.0.0.3/24**

3. Configure an alias name for the Gigabit Ethernet interface on the controller and configure its parameters.

   [edit]
   user@host# **set interfaces ge-0/0/0 vlan-tagging**
   user@host# **set interfaces ge-0/0/0 unit 0 alias "ge-to-corp-gw1"**
   user@host# **set interfaces ge-0/0/0.0 vlan-id 101**
   user@host# **set interfaces ge-0/0/0.0 family inet address 1.1.1.1/23**

4. Configure Gigabit Ethernet interfaces to be member links of an **ae-** logical interface.

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

5. Configure RIP in the network between the controller and the firewall gateway.

   [edit]
   user@host# **set protocols rip group corporate-firewall neighbor ge-to-corp-gw1**

Results

In configuration mode, confirm your configuration by entering the **show** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
        interfaces {
            ae0 {
                unit 0 {
                    alias "controller-sat1-downlink1";
                    family inet {
                        address 10.0.0.1/24;
                    }
                }
                unit 1 {
                    alias "controller-sat1-downlink2";
                    family inet {
                        address 10.0.0.3/24;
                    }
                }
            }
```

```
            ge-0/0/0 {
                vlan-tagging;
                unit 0 {
                    alias "ge-to-corp-gw1";
                    vlan-id 101;
                    family inet {
                        address 1.1.1.1/23;
                    }
                }
            }
            ge-0/1/0 {
  gigether-options {
                    802.3ad ae0;
                }
            }
            ge-0/1/1 {
  gigether-options {
                    802.3ad ae0;
                }
            }
        }
        protocols rip {
            group corporate-firewall {
                neighbor ge-to-corp-gw1;
            }
        }
```

After you have confirmed that the interfaces are configured, enter the **commit** command in configuration mode.

NOTE: In Juniper Networks M Series Multiservice Edge Routers, if the same alias name is configured on more than one logical interface, the router displays an error message and commit fails.

Verification

To verify that the alias name is displayed instead of the interface name, perform these steps:

- Verifying the Configuration of the Alias Name for the Controller Interfaces on page 126

Verifying the Configuration of the Alias Name for the Controller Interfaces

Purpose Verify that the alias name is displayed instead of the interface name.

 Copyright © 2018, Juniper Networks, Inc.

Action Display information about all RIP neighbors.

```
user@router> show rip neighbor
                  Local  Source          Destination     Send   Receive   In
Neighbor          State  Address         Address         Mode   Mode     Met
ge-to-corp-gw1       DN  (null)          255.255.255.255 mcast  both       1
```

Meaning The output displays the details of the benchmarking test that was performed. For more information about the **show rip neighbor** operational command, see **show rip neighbor** in the CLI Explorer.

Related Documentation

- Interface Alias Names Overview on page 122
- alias on page 435

Clock Source Overview

For both the router and interfaces, the clock source can be an external clock that is received on the interface or the router's internal Stratum 3 clock.

For example, interface A can transmit on interface A's received clock (external, loop timing) or the Stratum 3 clock (internal, line timing, or normal timing). Interface A cannot use a clock from any other source. For interfaces such as SONET/SDH that can use different clock sources, you can configure the source of the transmit clock on each interface.

The clock source resides on the System Control Board (SCB) for M40 routers, the System and Switch Board (SSB) for M20 routers, the Control Board (CB) for M120 routers, and the Miscellaneous Control Subsystem (MCS) for M40e and M160 routers. M7i and M10i routers have a clock source on the Compact Forwarding Engine Board (CFEB) and Enhanced Compact Forwarding Engine Board (CFEB-E).

For T Series and MX Series, the clock source internal Stratum 3 clock resides on the SONET Clock Generator and Switch Control Board (SCB) respectively. By default, the 19.44-MHz Stratum 3 reference clock generates the clock signal for all serial PICs (SONET/SDH) and Plesiochronous Digital Hierarchy (PDH) PICs. PDH PICs include DS3, E3, T1, and E1 PICs.

NOTE: **M7i and M10i routers do not support external clocking of SONET interfaces.**

For information about clocking on channelized interfaces, see *Channelized IQ and IQE Interfaces Properties*. Also see *Configuring the Clock Source on SONET/SDH Interfaces* and *Configuring the Channelized T3 Loop Timing*.

For information about configuring an external synchronization interface that can be used to synchronize the internal Stratum 3 clock to an external source on the M40e, M120,

Copyright © 2018, Juniper Networks, Inc.

M320, routers and T Series routers, see *Junos OS Administration Library*, *Configuring Junos OS to Support an External Clock Synchronization Interface for M Series, MX Series, and T Series Routers*.

For information about configuring Synchronous Ethernet on MX 80, MX240, MX480, and MX960 3D Universal Edge Routers, see *Junos OS Administration Library*, *Synchronous Ethernet Overview* and *Configuring Clock Synchronization Interface on MX Series Routers*.

Related Documentation

- *Configuring an External Synchronization Interface*
- Configuring the Clock Source on page 128
- *Configuring Junos OS to Support an External Clock Synchronization Interface for M Series, MX Series, and T Series Routers*
- *Synchronous Ethernet Overview*
- *Configuring Clock Synchronization Interface on MX Series Routers*

Configuring the Clock Source

For both the router and interfaces, the clock source can be an external clock that is received on the interface or the router's internal Stratum 3 clock.

To set the clock source as external or internal:

1. In configuration mode, go to the **[edit interfaces *interface-name*]** hierarchy level:

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

2. Configure the **clocking** option as external or internal.

   [edit interfaces *interface-name*]
   user@host# **set clocking (external | internal)**

NOTE: M7i and M10i routers do not support external clocking of SONET interfaces.

NOTE: On Channelized SONET/SDH PICs, if you set the parent (or the master) controller clock to **external**, then you must set the child controller clocks to the default value—that is, **internal**.

For example, on the Channelized STM1 PIC, if the clock on the Channelized STM1 interface (which is the master controller) is set to **external**, then you must not configure the CE1 interface (which is the child controller) clock to **external**. Instead you must configure the CE1 interface clock to **internal**.

 Copyright © 2018, Juniper Networks, Inc.

For information about clocking on channelized interfaces, see *Channelized IQ and IQE Interfaces Properties*. Also see *Configuring the Clock Source on SONET/SDH Interfaces* and *Configuring the Channelized T3 Loop Timing*.

For information about configuring an external synchronization interface that can be used to synchronize the internal Stratum 3 clock to an external source on the M40e, M120, and M320 routers and on the T Series routers, see *Junos OS Administration Library*, *Configuring Junos OS to Support an External Clock Synchronization Interface for M Series, MX Series, and T Series Routers*.

For information about configuring Synchronous Ethernet on MX80, MX240, MX480, and MX960 3D Universal Edge Routers, see *Junos OS Administration Library*, *Synchronous Ethernet Overview* and *Configuring Clock Synchronization Interface on MX Series Routers*.

Related Documentation

- *Configuring an External Synchronization Interface*
- **clocking on page 486**
- Clock Source Overview on page 127
- *Configuring Junos OS to Support an External Clock Synchronization Interface for M Series, MX Series, and T Series Routers*
- *Synchronous Ethernet Overview*
- *Configuring Clock Synchronization Interface on MX Series Routers*

Configuring Interface Encapsulation on Physical Interfaces

Understanding Interface Encapsulation on Physical Interfaces

Point-to-Point Protocol (PPP) encapsulation is the default encapsulation type for physical interfaces. You need not configure encapsulation for any physical interfaces that support PPP encapsulation. If you do not configure encapsulation, PPP is used by default.

For physical interfaces that do not support PPP encapsulation, you must configure an encapsulation to use for packets transmitted on the interface. You can optionally configure an encapsulation on a logical interface, which is the encapsulation used within certain packet types.

Encapsulation Capabilities of Physical Interfaces

When you configure a point-to-point encapsulation (such as PPP or Cisco HDLC) on a physical interface, the physical interface can have only one logical interface (that is, only one **unit** statement) associated with it. When you configure a multipoint encapsulation (such as Frame Relay), the physical interface can have multiple logical units, and the units can be either point-to-point or multipoint.

Copyright © 2018, Juniper Networks, Inc.

Ethernet CCC encapsulation for Ethernet interfaces with standard TPID tagging requires that the physical interface have only a single logical interface. Ethernet interfaces in VLAN mode can have multiple logical interfaces.

For Ethernet interfaces in VLAN mode, VLAN IDs are applicable as follows:

- VLAN ID 0 is reserved for tagging the priority of frames.
- For encapsulation type **vlan-ccc**, VLAN IDs 1 through 511 are reserved for normal VLANs. VLAN IDs 512 and above are reserved for VLAN CCCs.
- For encapsulation type **vlan-vpls**, VLAN IDs 1 through 511 are reserved for normal VLANs, and VLAN IDs 512 through 4094 are reserved for VPLS VLANs. For 4-port Fast Ethernet interfaces, you can use VLAN IDs 512 through 1024 for VPLS VLANs.
- For Gigabit Ethernet interfaces and Gigabit Ethernet IQ and IQE PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), you can configure flexible Ethernet services encapsulation on the physical interface. For interfaces with **flexible-ethernet-services** encapsulation, all VLAN IDs are valid. VLAN IDs from 1 through 511 are not reserved.
- For encapsulation types **extended-vlan-ccc** and **extended-vlan-vpls**, all VLAN IDs are valid.

The upper limits for configurable VLAN IDs vary by interface type.

When you configure a TCC encapsulation, some modifications are needed to handle VPN connections over unlike Layer 2 and Layer 2.5 links and terminate the Layer 2 and Layer 2.5 protocol locally.

The router performs the following media-specific changes:

- PPP TCC—Both Link Control Protocol (LCP) and Network Control Protocol (NCP) are terminated on the router. Internet Protocol Control Protocol (IPCP) IP address negotiation is not supported. The Junos OS strips all PPP encapsulation data from incoming frames before forwarding them. For output, the next hop is changed to PPP encapsulation.
- Cisco HDLC TCC—Keepalive processing is terminated on the router. The Junos OS strips all Cisco HDLC encapsulation data from incoming frames before forwarding them. For output, the next hop is changed to Cisco HDLC encapsulation.
- Frame Relay TCC—All Local Management Interface (LMI) processing is terminated on the router. The Junos OS strips all Frame Relay encapsulation data from incoming frames before forwarding them. For output, the next hop is changed to Frame Relay encapsulation.
- ATM—Operation, Administration, and Maintenance (OAM) and Interim Local Management Interface (ILMI) processing is terminated at the router. Cell relay is not supported. The Junos OS strips all ATM encapsulation data from incoming frames before forwarding them. For output, the next hop is changed to ATM encapsulation.

 Copyright © 2018, Juniper Networks, Inc.

Configuring the Encapsulation on a Physical Interface

By default, PPP is the encapsulation type for physical interfaces. To configure the encapsulation on a physical interface, include the encapsulation statement at the **[edit interfaces *interface-name*]** hierarchy level:

To configure encapsulation on a physical interface:

1. In configuration mode, go to the **[edit interfaces *interface-name*]** hierarchy level.

   [edit]
   user@host# **set interfaces so-*fpc/pic/port***

2. Configure the encapsulation type as described in encapsulation.

   [edit interfaces mo-*fpc/pic/port*]
   user@host# **set encapsulation *encapsulation-type***

NOTE:

- When you configure a point-to-point encapsulation (such as PPP or Cisco HDLC) on a physical interface, the physical interface can have only one logical interface (that is, only one unit statement) associated with it. When you configure a multipoint encapsulation (such as Frame Relay), the physical interface can have multiple logical units, and the units can be either point-to-point or multipoint.
- When the encapsulation type is set to Cisco-compatible Frame Relay encapsulation, ensure that the LMI type is set to ANSI or Q933-A.
- When vlan-vpls encapsulation is set at the physical interface level, commit check will validate that there should not be any inet family configured within it.

Displaying the Encapsulation on a Physical SONET/SDH Interface

Purpose To display the configured encapsulation and its associated set options on a physical interface when the following are set at the **[edit interfaces *interface-name*]** hierarchy level:

- interface-name—so-7/0/0
- Encapsulation—**ppp**
- Unit—0
- Family—**inet**
- Address—192.168.1.113/32
- Destination—192.168.1.114
- Family—**iso** and **mpls**

Copyright © 2018, Juniper Networks, Inc.

Action Run the **show** command at the **[edit interfaces *interface-name*]** hierarchy level.

```
[edit interfaces so-7/0/0]
user@host# show
encapsulation ppp;
unit 0 {
   point-to-point;
   family inet {
      address 192.168.1.113/32 {
         destination 192.168.1.114;
      }
   }
   family iso;
   family mpls;
}
```

Meaning The configured encapsulation and its associated set options are displayed as expected. Note that the second set of two **family** statements allow IS-IS and MPLS to run on the interface.

Related Documentation

- **encapsulation on page 577**
- Configuring the Media MTU on page 115

Configuring Interface Encapsulation on PTX Series Packet Transport Routers

This topic describes how to configure interface encapsulation on PTX Series Packet Transport Routers. Use the **flexible-ethernet-services** configuration statement to configure different encapsulation for different logical interfaces under a physical interface. With flexible Ethernet services encapsulation, you can configure each logical interface encapsulation without range restrictions for VLAN IDs.

Supported encapsulations for physical interfaces include:

- **flexible-ethernet-services**
- **ethernet-ccc**
- **ethernet-tcc**

Supported encapsulations for logical interfaces include:

- **ethernet**
- **vlan-ccc**
- **vlan-tcc**

 Copyright © 2018, Juniper Networks, Inc.

NOTE: PTX Series Packet Transport Routers do not support **extended-vlan-cc** and **extended-vlan-tcc** encapsulation on logical interfaces. Instead, you can configure a tag protocol ID (TPID) value of 0x9100 to achieve the same results.

To configure flexible Ethernet services encapsulation, include the **encapsulation flexible-ethernet-services** statement at the **[edit interfaces et-*fpc/pic/port*]** hierarchy level. For example:

```
interfaces {
  et-fpc/pic/port {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
      vlan-id 1000;
      family inet {
        address 11.0.0.20/24;
      }
    }
    unit 1 {
      encapsulation vlan-ccc;
      vlan-id 1010;
    }
    unit 2 {
      encapsulation vlan-tcc;
      vlan-id 1020;
      family tcc {
        proxy {
          inet-address 11.0.2.160;
        }
        remote {
          inet-address 11.0.2.10;
        }
      }
    }
  }
}
```

Related Documentation

- Configuring Interface Encapsulation on Physical Interfaces on page 129

Copyright © 2018, Juniper Networks, Inc.

Configuring Keepalives

By default, physical interfaces configured with Cisco HDLC or PPP encapsulation send keepalive packets at 10-second intervals. The Frame Relay term for keepalives is LMI packets; the Junos OS supports both ANSI T1.617 Annex D LMIs and ITU Q933 Annex A LMIs. On ATM networks, OAM cells perform the same function. You configure OAM cells at the logical interface level; for more information, see Defining the ATM OAM F5 Loopback Cell Period.

To disable the sending of keepalives:

1. In configuration mode, go to the **[edit interfaces *interface-name*]** hierarchy level.

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

2. Include the **no-keepalives** statement at the **[edit interfaces *interface-name*]** hierarchy level.

   ```
   [edit interfaces interface-name]
   no-keepalives;
   ```

To disable the sending of keepalives on a physical interface configured with Cisco HDLC encapsulation for a translational cross-connection:

1. In configuration mode, go to the **[edit interfaces*interface-name*]** hierarchy level.

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

2. Include the **no-keepalives** statement with the **encapsulation cisco-hdlc-tcc** statement at the **[edit interfaces *interface-name*]** hierarchy level.

   ```
   [edit interfaces interface-name]
   encapsulation cisco-hdlc-tcc;
   no-keepalives;
   ```

To disable the sending of keepalives on a physical interface configured with PPP encapsulation for a translational cross-connection:

1. In configuration mode, go to the **[edit interfaces *interface-name*]** hierarchy level.

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

2. Include the **no-keepalives** statement with the **encapsulation ppp-tcc** statement at the **[edit interfaces *interface-name*]** hierarchy level.

   ```
   [edit interfaces interface-name]
   encapsulation ppp-tcc;
   no-keepalives;
   ```

 Copyright © 2018, Juniper Networks, Inc.

For more information about translation cross-connections, see Circuit and Translational Cross-Connects Overview.

When you configure PPP over ATM or Multilink PPP over ATM encapsulation, you can enable or disable keepalives on the logical interface. For more information, see Configuring PPP over ATM2 Encapsulation.

To explicitly enable the sending of keepalives:

1. In configuration mode, go to the **[edit interfaces *interface-name*]** hierarchy level.

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

2. Include the **keepalives** statement at the **[edit interfaces *interface-name*]** hierarchy level.

   [edit interfaces*interface-name*]
   keepalives;

To change one or more of the default keepalive values:

1. In configuration mode, go to the **[edit interfaces *interface-name*]** hierarchy level.

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

2. Include the **keepalives** statement with the appropriate option as **interval*seconds***, **down-count*number***, and the **up-count*number***..

   [edit interfaces *interface-name*]
   keepalives;
   keepalives <interval *seconds*> <down-count *number*> <up-count *number*>;

On interfaces configured with Cisco HDLC or PPP encapsulation, you can include the following three keepalive statements; note that Frame Relay encapsulation is not affected by these statements:

- **interval *seconds***—The time in seconds between successive keepalive requests. The range is from 1 second through 32767 seconds, with a default of 10 seconds.
- **down-count *number***—The number of keepalive packets a destination must fail to receive before the network takes a link down. The range is from 1 through 255, with a default of 3.
- **up-count *number***—The number of keepalive packets a destination must receive to change a link's status from down to up. The range is from 1 through 255, with a default of 1.

CAUTION: If interface keepalives are configured on an interface that does not support the **keepalives** configuration statement (for example, 10-Gigabit Ethernet), the link layer may go down when the PIC is restarted. Avoid configuring the keepalives on interfaces that do not support the keepalives configuration statement.

For information about Frame Relay keepalive settings, see Configuring Frame Relay Keepalives.

On MX Series routers with Modular Port Concentrators/Modular Interface Cards (MPCs/MICs), the Packet Forwarding Engine on an MPC/MIC processes and responds to Link Control Protocol (LCP) Echo-Request keepalive packets that the PPP subscriber (client) initiates and sends to the router. The mechanism by which LCP Echo-Request packets are processed by the Packet Forwarding Engine instead of by the Routing Engine is referred to as *PPP fast keepalive* For more information about how PPP fast keepalive works on an MX Series router with MPCs/MICs, see the *Junos OS Subscriber Access Configuration Guide*.

Related Documentation

- *Defining the ATM OAM F5 Loopback Cell Period*
- *Disabling the Sending of PPPoE Keepalive Messages*
- *Understanding How the Router Processes Subscriber-Initiated PPP Fast Keepalive Requests*
- **keepalives on page 710**
- **no-keepalives on page 817**
- *Configuring Frame Relay Keepalives*
- Circuit and Translational Cross-Connects Overview on page 275
- *Configuring PPP over ATM2 Encapsulation Overview*

Configuring the PPP Challenge Handshake Authentication Protocol

PPP Challenge Handshake Authentication Protocol

For interfaces with PPP encapsulation, you can configure interfaces to support the PPP Challenge Handshake Authentication Protocol (CHAP), as defined in RFC 1994, PPP Challenge Handshake Authentication Protocol (CHAP). When you enable CHAP on an interface, the interface can authenticate its peer and can be authenticated by its peer. By default, PPP CHAP is disabled. If CHAP is not explicitly enabled, the interface makes no CHAP challenges and denies all incoming CHAP challenges. To enable CHAP, you must create an access profile, and you must configure the interfaces to use CHAP.

Copyright © 2018, Juniper Networks, Inc.

Configuring the PPP Challenge Handshake Authentication Protocol

When you configure an interface to use CHAP, you must assign an access profile to the interface. When an interface receives CHAP challenges and responses, the access profile in the packet is used to look up the shared secret, as defined in RFC 1994. If no matching access profile is found for the CHAP challenge that was received by the interface, the optionally configured default CHAP secret is used. The default CHAP secret is useful if the CHAP name of the peer is unknown, or if the CHAP name changes during PPP link negotiation.

To enable CHAP, you must create an access profile, and you must configure the interfaces to use PAP. For more information on how to configure access profile, see *Configuring Access Profiles for L2TP or PPP Parameters*.

To configure the PPP challenge handshake authentication protocol, on each physical interface with PPP encapsulation, perform the following steps.

1. To assign an access profile to an interface, include the **access-profile** statement at the **[edit interfaces interface-name ppp-options chap]** hierarchy level.

   [edit interfaces *interface-name* ppp-options chap]
   user@host# **set** access-profile ***name***

   NOTE: You must include the **access-profile** statement when you configure the CHAP authentication method. If an interface receives a CHAP challenge or response from a peer that is not in the applied access profile, the link is immediately dropped unless a default CHAP secret has been configured.

2. The default CHAP secret is used when no matching CHAP access profile exists, or if the CHAP name changes during PPP link negotiation. To configure a default CHAP secret for an interface, include the **default-chap-secret** statement at the **[edit interfaces *interface-name* ppp-options chap]** hierarchy level.

   [edit interfaces *interface-name* ppp-options chap]
   user@host# **set** default-chap-secret ***name***

3. To configure the name the interface uses in CHAP challenge and response packets, include the **local-name** statement at the **[edit interfaces *interface-name* ppp-options chap]** hierarchy level:

   [edit interfaces *interface-name* ppp-options chap]
   user@host# **set** local-name ***name***

Copyright © 2018, Juniper Networks, Inc.

NOTE:

- The local name is any string from 1 through 32 characters in length, starting with an alphanumeric or underscore character, and including only the following characters:

  a-z A-Z 0-9 % @ # / \ . _ -

- By default, when CHAP is enabled on an interface, the interface uses the router's system hostname as the name sent in CHAP challenge and response packets.

4. You can configure the interface not to challenge its peer, and only respond when challenged. To configure the interface not to challenge its peer, include the **passive** statement at the **[edit interfaces *interface-name* ppp-options chap]** hierarchy level:

```
[edit interfaces interface-name ppp-options chap]
user@host# set passive;
```

NOTE: By default, when CHAP is enabled on an interface, the interface always challenges its peer and responds to challenges from its peer.

See Also

- Configuring the PPP Authentication Protocol on page 142

Displaying the Configured PPP Challenge Handshake Authentication Protocol

Purpose To display the configured PPP CHAP at the **[edit access]** and **[edit interfaces]** hierarchy levels.

- Access profile—**pe-A-ppp-clients**
- default CHAP secret data—**"$ABC123"**
- hostname for the CHAP challenge and response packets—**"pe-A-so-1/1/1"**
- Interface—so-1/1/2

Action

- Run the **show** command at the **[edit access]** hierarchy level.

```
profile pe-A-ppp-clients;
client cpe-1 chap-secret "$ABC123";
        # SECRET-DATA
[edit interfaces so-1/2/0]
encapsulation ppp;
ppp-options {
  chap {
    access-profile pe-A-ppp-clients;
    default-chap-secret "$ABC123";
    local-name "pe-A-so-1/1/1";
  }
```


Copyright © 2018, Juniper Networks, Inc.

```
}
```

- Run the **show** command at the **[edit interfaces s0-1/1/2]** hierarchy level.

```
ppp-options {
  chap {
    access-profile pe-A-ppp-clients;
    default-chap-secret "$ABC123";
    local-name “pe-A-so-1/1/2";
  }
}
```

Meaning The configured CHAP and its associated set options are displayed as expected.

Configuring the PPP Password Authentication Protocol On a Physical Interface

Understanding PPP Password Authentication Protocol

For interfaces with PPP encapsulation, you can configure interfaces to support the Password Authentication Protocol (PAP), as defined in RFC 1334, *PAP Authentication Protocols*. If authentication is configured, the PPP link negotiates using CHAP or PAP protocol for authentication during the Link Control Protocol (LCP) negotiation phase. PAP is only performed after the link establishment phase (LCP up) portion of the authentication phase.

During authentication, the PPP link sends a PAP authentication-request packet to the peer with an ID and password. The authentication-request packet is sent every 2 seconds, similar to the CHAP challenge, until a response is received (acknowledgment packet, nonacknowledgment packet). If an acknowledgment packet is received, the PPP link transitions to the next state, the network phase. If a nonacknowledgment packet is received, an LCP terminate request is sent, and the PPP link goes back to the link establishment phase. If no response is received, and an optional retry counter is set to **true**, a new request acknowledgment packet is resent. If the retry counter expires, the PPP link transitions to the LCP negotiate phrase.

You can configure the PPP link with PAP in passive mode. By default, when PAP is enabled on an interface, the interface expects authenticate-request packets from the peer. However, the interface can be configured to send authentication request packets to the peer by configuring PAP to operate in passive mode. In PAP passive mode, the interface sends the authenticate-request packets to the peer only if the interface receives the PAP option from the peer during LCP negotiation—in passive mode, the interface does not authenticate the peer.

Copyright © 2018, Juniper Networks, Inc.

Configuring the PPP Password Authentication Protocol On a Physical Interface

When you configure an interface to use PAP, you must assign an access profile to the interface. When an interface receives PAP authentication requests, the access profile in the packet is used to look up the password.

To enable PAP, you must create an access profile, and you must configure the interfaces to use PAP. For more information on how to configure access profile, see *Configuring Access Profiles for L2TP or PPP Parameters*.

To configure the PPP password authentication protocol, on each physical interface with PPP encapsulation, perform the following steps.

1. To assign an access profile to an interface, include the **access-profile** statement at the **[edit interfaces interface-name ppp-options pap]** hierarchy level.

   [edit interfaces *interface-name* ppp-options pap]
   user@host# **set** access-profile***name***

2. To configure the name the interface uses in PAP request and response packets, include the **local-name** statement at the **[edit interfaces *interface-name* ppp-options pap]** hierarchy level:

   [edit interfaces *interface-name* ppp-options pap]
   user@host# **set** local-name ***name***

3. You need to configure the password to be used for authentication. To configure the host password for sending PAP requests, include the **local-password** statement at the **[edit interfaces *interface-name* ppp-options pap]** hierarchy level:

   [edit interfaces *interface-name* ppp-options pap]
   user@host# **set** local-password ***password***

   NOTE: By default, when PAP is enabled on an interface, the interface uses the router's system hostname as the name sent in PAP request and response packets.

4. To configure the interface to authenticate with PAP in passive mode, include the **passive** statement at the **[edit interfaces *interface-name* ppp-options pap]** hierarchy level:

   [edit interfaces *interface-name* ppp-options pap]
   user@host# **set** passive

 Copyright © 2018, Juniper Networks, Inc.

NOTE: By default, when PAP is enabled on an interface, the interface expects authenticate-request packets from the peer. However, the interface can be configured to send authentication request packets to the peer by configuring PAP to operate in passive mode. In PAP passive mode, the interface sends the authenticate-request packets to the peer only if the interface receives the PAP option from the peer during LCP negotiation—in passive mode, the interface does not authenticate the peer.

See Also
- Configuring the PPP Authentication Protocol on page 142

Configuring the PPP Password Authentication Protocol On a Logical Interface

When you configure an interface to use PAP, you must assign an access profile to the interface. When an interface receives PAP authentication requests, the access profile in the packet is used to look up the password. If no matching access profile is found for the PAP authentication request that was received by the interface, the optionally configured default PAP password is used.

To configure the PPP password authentication protocol, on each logical interface with PPP encapsulation, perform the following steps.

1. To configure the default PAP password, include the **pap-password** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* ppp-options pap]** hierarchy level:

   [edit interfaces *interface-name* unit *logical-unt-number* ppp-options pap]
   user@host# **set** default-pap-password ***password***

2. To configure the name the interface uses in PAP request and response packets, include the **local-name** statement at the **[edit interfaces *interface-name* unit *logical-unt-number* ppp-options pap]** hierarchy level:

   [edit interfaces *interface-name* ppp-options pap]
   user@host# **set** local-name ***name***

3. You need to configure the password to be used for authentication. To configure the host password for sending PAP requests, include the **local-password** statement at the **[edit interfaces *interface-name* ppp-options pap]** hierarchy level:

   [edit interfaces *interface-name* unit *logical-unt-number* ppp-options pap]
   user@host# **set** local-password ***password***

NOTE: By default, when PAP is enabled on an interface, the interface uses the router's system hostname as the name sent in PAP request and response packets.

4. To configure the interface to authenticate with PAP in passive mode, include the **passive** statement at the **[edit interfaces *interface-name* unit *logical-unt-number*ppp-options pap]** hierarchy level:

   [edit interfaces *interface-name* unit *logical-unt-number* ppp-options pap]
   user@host# **set** passive

   NOTE: By default, when PAP is enabled on an interface, the interface expects authenticate-request packets from the peer. However, the interface can be configured to send authentication request packets to the peer by configuring PAP to operate in passive mode. In PAP passive mode, the interface sends the authenticate-request packets to the peer only if the interface receives the PAP option from the peer during LCP negotiation—in passive mode, the interface does not authenticate the peer.

See Also

• Configuring the PPP Authentication Protocol on page 142

Configuring the PPP Authentication Protocol

The Point-to-Point Protocol (PPP) is an encapsulation protocol for transporting IP traffic across point-to-point links. To configure PPP, you can configure the Challenge Handshake Authentication Protocol (CHAP). CHAP allows each end of a PPP link to authenticate its peer, as defined in RFC 1994. The authenticator sends its peer a randomly-generated challenge that the peer must encrypt using a one-way hash; the peer must then respond with that encrypted result. The key to the hash is a secret known only to the authenticator and authenticated. When the response is received, the authenticator compares its calculated result with the peer's response. If they match, the peer is authenticated.

Each end of the link identifies itself to its peer by including its name in the CHAP challenge and response packets it sends to the peer. This name defaults to the local hostname, or you can explicitly set it using the **local-name** option. When a host receives a CHAP challenge or CHAP response packet on a particular interface, it uses the peer identity to look up the CHAP secret key to use.

To configure CHAP, include the **profile** statement at the **[edit access]** hierarchy level:

```
[edit access]
profile profile-name {
  client client-name chap-secret chap-secret;
}
```

Then reference the CHAP profile name at the **[edit interfaces]** hierarchy level.

You can configure multiple CHAP profiles, and configure multiple clients for each profile.

 Copyright © 2018, Juniper Networks, Inc.

Definitions:

- **profile** is the mapping between peer identifiers and CHAP secret keys. The identity of the peer contained in the CHAP challenge or response queries the profile for the secret key to use.
- **client** is the peer identity.
- **chap-secret** is the secret key associated with that peer.

Related Documentation

- *Example: Configuring PPP CHAP*
- *Example: Configuring CHAP Authentication with RADIUS*

PPP Encapsulation on ACX Series Routers

You can configure Point-to-Point Protocol (PPP) encapsulation on physical interfaces on ACX Series routers. PPP provides a standard method for transporting multiprotocol datagrams over a point-to-point link. PPP uses the High-Speed Data Link Control (HDLC) protocol for its physical interface and provides a packet- oriented interface for the network-layer protocols.

PPP is supported on the following MICs on ACX Series routers:

- On ACX1000 routers with 8-port built-in T1/E1 TDM MICs.
- On ACX2000, ACX2100, ACX2200, and ACX4000 routers with 16-port built-in T1/E1 TDM MICs.
- On ACX4000 routers with 16-Port Channelized E1/T1 Circuit Emulation MICs.

Starting with Release 12.3X54, you can configure Point-to-Point Protocol (PPP) encapsulation on physical interfaces on Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP on ACX4000 Series routers

On ACX Series routers, E1, T1, and *N*xDS0 interfaces support PPP encapsulation.

PPP is the default encapsulation type for physical interfaces. You need not configure encapsulation for any physical interfaces that support PPP encapsulation. If you do not configure encapsulation, PPP is used by default. For physical interfaces that do not support PPP encapsulation, you must configure an encapsulation to use for packets transmitted on the interface.

To configure the encapsulation on a physical interface, include the **encapsulation ppp** statement at the **[edit interfaces *interface-name*]** hierarchy level.

IP class of service (CoS) is not supported on PPP interfaces. All the traffic is sent to the best effort queue (queue 0) and CoS code points are not processed. Also, fixed classifiers are not supported. Circuit cross-connect (CCC) version of PPP (**ppp-ccc** option) and translational cross-connect (TCC) version of PPP (**ppp-tcc** option) are not supported for configuration with the **encapsulation** statement.

Copyright © 2018, Juniper Networks, Inc.

PPP is supported only for IPv4 networks. If you configure PPP encapsulation, you can configure an INET family by including the **family inet** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level. MPLS family is not supported on logical interfaces if you configured PPP encapsulation. On interfaces with PPP encapsulation, configure PPP-specific interface properties by including the **ppp-options** statement at the **[edit interfaces *interface-name*]** hierarchy level. For interfaces with PPP encapsulation, you can configure interfaces to support the PPP Challenge Handshake Authentication Protocol (CHAP) and Password Authentication Protocol (PAP).

For full T1/E1 interfaces on which PPP encapsulation needs to be enabled, create the T1/E1 interfaces out of channelized T1/E1 interfaces (CT1/CE1) by including the **framing** statement at the **[edit chassis fpc *fpc-slot* pic *pic-slot*]** hierarchy level:

```
[edit chassis fpc fpc-slot pic pic-slot]
user@host# set framing (t1 | e1);
```

Configure a CT1 port down to a T1 channel. On the CT1 interface, set the **no-partition** option and then set the interface type as T1.

```
[edit interfaces ct1-mpc-slot/mic-slot/port-number]
user@host# set no-partition interface-type t1
```

Configure a CE1 port down to an E1 channel. On the CE1 interface, set the **no-partition** option and then set the interface type as E1.

```
[edit interfaces ce1-mpc-slot/mic-slot/port-number]
user@host# set no-partition interface-type t1
```

For *N*xDS0 interfaces on which PPP encapsulation needs to be enabled, partition the CE1 and CT1 interfaces by including the **ce1-*x/y/z* partition *partition-number* timeslots *timeslots* interface-type ds** and **ct1-*x/y/z* partition *partition-number* timeslots *timeslots* interface-type ds** statements at the [**edit interfaces *interface-name***] hierarchy level.

The following operational mode commands can be used to view PPP configuration settings and statistical details:

- The **show ppp address-pool** command is used to display PPP address pool information.
- The **show ppp interface** command is used to display PPP session information for an interface.
- The **show ppp statistics** command is used to display PPP session statistics.
- The **show ppp summary** command is used to display summary information about PPP-configured interfaces.
- The **show interfaces e1-*fpc/pic/port***, **show interfaces t1-*fpc/pic/port***, and **show interfaces ds-*fpc/pic/port*** commands are used to display the PPP settings of a specific E1, T1, and DS interface, respectively.

Related Documentation

- Configuring Interface Encapsulation on Physical Interfaces in ACX Series on page 145
- *encapsulation*
- **ppp-options on page 891**

 Copyright © 2018, Juniper Networks, Inc.

Configuring Interface Encapsulation on Physical Interfaces in ACX Series

Point-to-Point Protocol (PPP) encapsulation is the default encapsulation type for physical interfaces. You need not configure encapsulation for any physical interfaces that support PPP encapsulation. If you do not configure encapsulation, PPP is used by default. For physical interfaces that do not support PPP encapsulation, you must configure an encapsulation to use for packets transmitted on the interface.

You can optionally configure an encapsulation on a logical interface, which is the encapsulation used within certain packet types. For more information about logical interface encapsulation, see *Configuring Interface Encapsulation on Logical Interfaces*.

This section contains the following topics:

Configuring the Encapsulation on a Physical Interface

By default, PPP is the encapsulation type for physical interfaces. To configure the encapsulation on a physical interface, include the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
encapsulation (atm-ccc-cell-relay | atm-pvc | cisco-hdlc | cisco-hdlc-ccc | cisco-hdlc-tcc
  | ethernet-ccc | ethernet-over-atm | ethernet-tcc | ethernet-vpls |
  extended-frame-relay-ccc | extended-frame-relay-ether-type-tcc |
  extended-frame-relay-tcc | extended-vlan-ccc | extended-vlan-tcc | extended-vlan-vpls
  | flexible-ethernet-services | flexible-frame-relay | frame-relay | frame-relay-ccc |
  frame-relay-ether-type | frame-relay-ether-type-tcc | frame-relay-port-ccc |
  frame-relay-tcc | multilink-frame-relay-uni-nni | ppp | ppp-ccc | ppp-tcc | vlan-ccc |
  vlan-vpls);
```

NOTE: ACX Series routers do not support *cisco-hdlc* encapsulation.

The physical interface encapsulation can be one of the following:

- ATM CCC cell relay—Connects two remote virtual circuits or ATM physical interfaces with a label-switched path (LSP). Traffic on the circuit is ATM cells.

  For more information, see the *Junos OS Administration Library*.

- ATM PVC—Defined in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*. When you configure physical ATM interfaces with ATM PVC encapsulation, an RFC 2684-compliant ATM Adaptation Layer 5 (AAL5) tunnel is set up to route the ATM cells over a Multiprotocol Label Switching (MPLS) path that is typically established between two MPLS-capable routers using the Label Distribution Protocol (LDP).

- Ethernet cross-connect—Ethernet interfaces without VLAN tagging can use Ethernet CCC encapsulation. Two related versions are supported:

- CCC version (**ethernet-ccc**)—Ethernet interfaces with standard Tag Protocol ID (TPID) tagging can use Ethernet CCC encapsulation. When you use this encapsulation type, you can configure the **ccc** family only.
- TCC version (**ethernet-tcc**)—Similar to CCC, but used for circuits with different media on either side of the connection.

  For 8-port, 12-port, and 48-port Fast Ethernet PICs, TCC is not supported.

- VLAN CCC (**vlan-ccc**)—Ethernet interfaces with VLAN tagging enabled can use VLAN CCC encapsulation. VLAN CCC encapsulation supports TPID 0x8100 only. When you use this encapsulation type, you can configure the **ccc** family only.
- Extended VLAN cross-connect—Gigabit Ethernet interfaces with VLAN 802.1Q tagging enabled can use extended VLAN cross-connect encapsulation. (Ethernet interfaces with standard TPID tagging can use VLAN CCC encapsulation.) Two related versions of extended VLAN cross-connect are supported:
  - CCC version (**extended-vlan-ccc**)—Extended VLAN CCC encapsulation supports TPIDs 0x8100, 0x9100, and 0x9901. When you use this encapsulation type, you can configure the **ccc** family only.
  - TCC version (**extended-vlan-tcc**)—Similar to CCC, but used for circuits with different media on either side of the connection.

    For 8-port, 12-port, and 48-port Fast Ethernet PICs, extended VLAN CCC is not supported. For 4-port Gigabit Ethernet PICs, extended VLAN CCC and extended VLAN TCC are not supported.

NOTE: In ACX Series routers, VPLS is supported only on ACX5048 and ACX5096 routers.

- Ethernet VPLS (**ethernet-vpls**)—Ethernet interfaces with VPLS enabled can use Ethernet VPLS encapsulation. For more information about VPLS, see the *Junos OS VPNs Library for Routing Devices*.
- Ethernet VLAN VPLS (**vlan-vpls**)—Ethernet interfaces with VLAN tagging and VPLS enabled can use Ethernet VLAN VPLS encapsulation. For more information about VPLS, see the *Junos OS VPNs Library for Routing Devices*.
- Extended VLAN VPLS (**extended-vlan-vpls**)—Ethernet interfaces with VLAN 802.1Q tagging and VPLS enabled can use Ethernet Extended VLAN VPLS encapsulation. (Ethernet interfaces with standard TPID tagging can use Ethernet VLAN VPLS encapsulation.) Extended Ethernet VLAN VPLS encapsulation supports TPIDs 0x8100, 0x9100, and 0x9901. For more information about VPLS, see the *Junos OS VPNs Library for Routing Devices*.
- Flexible Ethernet services (**flexible-ethernet-services**)—Gigabit Ethernet and Gigabit Ethernet IQ and IQE PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router) can use flexible Ethernet services encapsulation. Aggregated Ethernet bundles can use this encapsulation type. You use this encapsulation type when you want to configure multiple per-unit Ethernet

 Copyright © 2018, Juniper Networks, Inc.

encapsulations. This encapsulation type allows you to configure any combination of route, TCC, CCC, Layer 2 virtual private networks (VPNs), and VPLS encapsulations on a single physical port. If you configure flexible Ethernet services encapsulation on the physical interface, VLAN IDs from 1 through 511 are no longer reserved for normal VLANs.

- PPP—Defined in RFC 1661, *The Point-to-Point Protocol (PPP) for the Transmission of Multiprotocol Datagrams over Point-to-Point Links*. PPP is the default encapsulation type for physical interfaces. E1, E3, SONET/SDH, T1, and T3 interfaces can use PPP encapsulation.

NOTE: When the encapsulation type is set to Cisco-compatible Frame Relay encapsulation, ensure that the LMI type is set to ANSI or Q933-A.

In ACX Series routers, VPLS is supported only on ACX5048 and ACX5096 routers.

Encapsulation Capabilities

When you configure a point-to-point encapsulation (such as PPP or Cisco HDLC) on a physical interface, the physical interface can have only one logical interface (that is, only one **unit** statement) associated with it. When you configure a multipoint encapsulation (such as Frame Relay), the physical interface can have multiple logical units, and the units can be either point-to-point or multipoint.

Ethernet CCC encapsulation for Ethernet interfaces with standard TPID tagging requires that the physical interface have only a single logical interface. Ethernet interfaces in VLAN mode can have multiple logical interfaces.

For Ethernet interfaces in VLAN mode, VLAN IDs are applicable as follows:

- VLAN ID 0 is reserved for tagging the priority of frames.
- For encapsulation type **vlan-ccc**, VLAN IDs 1 through 511 are reserved for normal VLANs. VLAN IDs 512 and above are reserved for VLAN CCCs.

  When you configure Ethernet virtual LAN (VLAN) encapsulation on CCC circuits (by using the **encapsulation vlan-ccc** statement at the **[edit interfaces *interface-name*]** hierarchy level), you can bind a list of VLAN IDs to the interface by using the **vlan-id-list [*vlan-id-numbers*]** statement to configure a CCC for multiple VLANs. Configuring this statement creates a CCC for:

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

- For encapsulation type **vlan-vpls**, VLAN IDs 1 through 511 are reserved for normal VLANs, and VLAN IDs 512 through 4094 are reserved for VPLS VLANs. For 4-port Fast Ethernet interfaces, you can use VLAN IDs 512 through 1024 for VPLS VLANs.

Copyright © 2018, Juniper Networks, Inc.

- For Gigabit Ethernet interfaces and Gigabit Ethernet IQ and IQE PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), you can configure flexible Ethernet services encapsulation on the physical interface. For interfaces with **flexible-ethernet-services** encapsulation, all VLAN IDs are valid. VLAN IDs from 1 through 511 are not reserved.
- For encapsulation types **extended-vlan-ccc** and **extended-vlan-vpls**, all VLAN IDs are valid.

The upper limits for configurable VLAN IDs vary by interface type.

When you configure a TCC encapsulation, some modifications are needed to handle VPN connections over unlike Layer 2 and Layer 2.5 links and terminate the Layer 2 and Layer 2.5 protocol locally.

The router performs the following media-specific change:

- ATM—Operation, Administration, and Maintenance (OAM) and Interim Local Management Interface (ILMI) processing is terminated at the router. Cell relay is not supported. The Junos OS strips all ATM encapsulation data from incoming frames before forwarding them. For output, the next hop is changed to ATM encapsulation.

Example: Configuring the Encapsulation on a Physical Interface

Configure PPP encapsulation on a SONET/SDH interface. The second and third **family** statements allow Intermediate System-to-Intermediate System (IS-IS) and MPLS to run on the interface.

```
[edit interfaces]
so-7/0/0 {
  encapsulation ppp;
  unit 0 {
    point-to-point;
    family inet {
      address 192.168.1.113/32 {
        destination 192.168.1.114;
      }
    }
    family iso;
    family mpls;
  }
}
```

Related Documentation

- *Configuring Interface Encapsulation on Logical Interfaces*

Copyright © 2018, Juniper Networks, Inc.

Configuring PPP Address and Control Field Compression

For interfaces with PPP, PPP CCC, or PPP TCC encapsulation, you can configure compression of the Data Link Layer address and control fields, as defined in RFC 1661, *The Point-to-Point Protocol (PPP)*. By default, the address and control fields are not compressed. This means PPP-encapsulated packets are transmitted with two 1-byte fields (0xff and 0x03). If you configure address and control field compression (ACFC) and ACFC is successfully negotiated with the local router's peer, the local router transmits packets without these 2 bytes. ACFC allows you to conserve bandwidth by transmitting less data.

On M320, M120, and T Series routers, ACFC is not supported for any ISO family protocols. Do not include the **acfc** statement at the **[edit interfaces *interface-name* ppp-options compression]** hierarchy level when you include the **family iso** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

NOTE: The address and control fields cannot be compressed in Link Control Protocol (LCP) packets.

The PPP session restarts when you configure or modify compression options.

To configure ACFC:

1. In configuration mode, go to the **[edit interfaces *interface-name* ppp-options]** hierarchy level.

   [edit]
   user@host# **edit interfaces *interface-name* ppp-options**

2. Include the **compression** statement at the **[edit interfaces *interface-name* ppp-options]** hierarchy level, and specify **acfc**.

   [edit interfaces *interface-name* ppp-options]
   compression acfc;

To monitor the configuration, issue the **show interfaces *interface-name*** command. Configured options are displayed in the **link flags** field for the physical interface. Successfully negotiated options are displayed in the **flags** field for the logical interface. In this example, both ACFC and PFC are configured, but neither compression feature has been successfully negotiated.

```
user@router# run show interfaces so-0/1/1
Physical interface: so-0/1/1, Enabled, Physical link is Up
 Interface index: 133, SNMP ifIndex: 27
 Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC3,
Loopback: None, FCS: 16
  Payload scrambler: Enabled
  Device flags    : Present Running
  Interface flags: Point-To-Point SNMP-Traps 16384
  Link flags      : No-Keepalives ACFC PFC
```

Copyright © 2018, Juniper Networks, Inc.

```
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
  CHAP state: Not-configured
  CoS queues     : 4 supported
  Last flapped   : 2004-12-29 10:49:32 PST (00:18:35 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : None
  SONET defects  : None
  Logical interface so-0/1/1.0 (Index 68) (SNMP ifIndex 169)
    Flags: Point-To-Point SNMP-Traps ACFC Encapsulation: PPP
    Protocol inet, MTU: 4470
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 3.3.3/24, Local: 3.3.3.2, Broadcast: 3.3.3.255
```

This configuration causes the local router to try to negotiate ACFC with its peer. If ACFC is successfully negotiated, the local router sends packets with compressed address and control fields. When you include the **compression acfc** statement in the configuration, the PPP session restarts, and the local router sends the ACFC option in the LCP Configure-Request packet. The ACFC option informs the local router's peer that the local router can receive packets with compression. If the peer indicates that it, too, can receive packets with compression, then ACFC is negotiated. If ACFC is successfully negotiated, the local router can receive packets with or without the address and control bytes included.

Related Documentation

- ppp-options on page 891
- compression on page 492
- acfc on page 415

 Copyright © 2018, Juniper Networks, Inc.

Configuring the PPP Protocol Field Compression

For interfaces with PPP, PPP CCC, or PPP TCC encapsulation, you can configure protocol field compression. By default, the protocol field is not compressed. This means PPP-encapsulated packets are transmitted with a two-byte protocol field. For example, IPv4 packets are transmitted with the protocol field set to 0x0021, and MPLS packets are transmitted with the protocol field set to 0x0281.

For all protocols with identifiers in the range 0x0000 through 0x00ff, you can configure the router to compress the protocol field to one byte, as defined in RFC 1661, *The Point-to-Point Protocol (PPP)*. Protocol field compression (PFC) allows you to conserve bandwidth by transmitting less data.

NOTE: The protocol field cannot be compressed in Link Control Protocol (LCP) packets.

The PPP session restarts when you configure or modify compression options.

To configure PFC:

1. In configuration mode, go to the **[edit interfaces *interface-name* ppp-options]** hierarchy level.

   [edit]
   user@host# **edit interfaces *interface-name* ppp-options**

2. Include the **compression** statement at the **[edit interfaces *interface-name* ppp-options]** hierarchy level, and specify **pfc**.

   [edit interfaces *interface-name* ppp-options]
   compression pfc;

To monitor the configuration, issue the **show interfaces *interface-name*** command. Configured options are displayed in the **link flags** field for the physical interface. Successfully negotiated options are displayed in the **flags** field for the logical interface. In this example, both ACFC and PFC are configured, but neither compression feature has been successfully negotiated.

```
user@router# run show interfaces so-0/1/1
Physical interface: so-0/1/1, Enabled, Physical link is Up
  Interface index: 133, SNMP ifIndex: 27
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC3,
Loopback: None, FCS: 16,
  Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 16384
  Link flags     : No-Keepalives ACFC PFC
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
  CHAP state: Not-configured
```

Copyright © 2018, Juniper Networks, Inc.

```
CoS queues     : 4 supported
Last flapped   : 2004-12-29 10:49:32 PST (00:18:35 ago)
Input rate     : 0 bps (0 pps)
Output rate    : 0 bps (0 pps)
SONET alarms   : None
SONET defects  : None
Logical interface so-0/1/1.0 (Index 68) (SNMP ifIndex 169)
  Flags: Point-To-Point SNMP-Traps ACFC Encapsulation: PPP
  Protocol inet, MTU: 4470
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 3.3.3/24, Local: 3.3.3.2, Broadcast: 3.3.3.255
```

This configuration causes the local router to try to negotiate PFC with its peer. If PFC is successfully negotiated, the local router sends packets with compressed protocol fields. When you include the **compression pfc** statement in the configuration, the PPP session restarts, and the local router sends the PFC option in the LCP Configure-Request packet. The PFC option informs the local router's peer that the local router can receive packets with compression. If the peer indicates that it, too, can receive packets with compression, then PFC is negotiated. If PFC is successfully negotiated, the local router can receive packets with either 2-byte (uncompressed) or 1-byte (compressed) protocol fields.

Related Documentation

ppp-options on page 891

compression on page 492

pfc on page 870

Monitoring a PPP Session

You can monitor PPP packet exchanges. When monitoring is enabled, packets exchanged during a session are logged by default to **/var/log/pppd**, or to the file specified in the **traceoptions** statement.

To monitor a PPP session:

1. In configuration mode, go to the **[edit protocols ppp]** hierarchy level.

   [edit]
   user@host# **edit protocols ppp**

2. Include the **monitor-session** statement.

   [edit protocols ppp]
   user@host# **monitor-session (*interface-name* | all);**

When monitoring is configured, the operational mode commands **show ppp summary** and **show ppp interface** display a **Monitored** flag in the **Session flags** column or line.

Related Documentation

monitor-session on page 782

 Copyright © 2018, Juniper Networks, Inc.

Tracing Operations of the pppd Process

You can trace the operations of the router's pppd process.

To trace the router's pppd process:

1. In configuration mode, go to the **[edit protocols ppp]** hierarchy level.

```
[edit ]
user@host# edit protocols ppp
```

2. Include the **traceoptions** statement.

```
[edit protocols ppp]
traceoptions {
  file filename <files number> <match regular-expression> <size size> <world-readable |
    no-world-readable>;
  flag flag;
  level severity-level;
  no-remote-trace;
}
```

- To specify more than one tracing operation, include multiple **flag** statements.

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

  - **access**—Trace access code
  - **address-pool**—Trace address pool code
  - **all**—Trace all areas of code
  - **auth**—Trace authentication code
  - **chap**—Trace challenge handshake authentication protocol code
  - **ci**—Trace CI code
  - **config**—Trace configuration code
  - **ifdb**—Trace interface database code
  - **lcp**—Trace LCP state machine code
  - **memory**—Trace memory management code
  - **message**—Trace message processing code
  - **mlppp**—Trace multilink point-to-point protocol code
  - **ncp**—Trace NCP state machine code
  - **pap**—Trace password authentication protocol code
  - **ppp**—Trace PPP protocol processing code
  - **radius**—Trace RADIUS processing code
  - **redundancy**—Trace redundancy code

Copyright © 2018, Juniper Networks, Inc.

- **rtsock**—Trace routing socket code
- **session**—Trace session management code
- **signal**—Trace signal handling code
- **timer**—Trace timer code
- **ui**—Trace user interface code

Related Documentation traceoptions on page 1033

Configuring the Router as a DCE with Frame Relay Encapsulation

By default, when you configure an interface with Frame Relay encapsulation, the routing platform is assumed to be data terminal equipment (DTE). That is, the routing platform is assumed to be at a terminal point on the network.

When you configure the router to be a data circuit-terminating equipment (DCE), keepalives are disabled by default.

To configure the router to be DCE:

1. In configuration mode, go to the **[edit interfaces *interface-name*]** hierarchy level:

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

2. Configure the **dce** option to respond to Frame Relay status enquiry messages.

   [edit interfaces *interface-name*]
   user@host# **set dce**

Perform one of the following tasks for back-to-back Frame Relay connections:

- Disable sending of keepalives on both sides of the connection.
- Configure one side of the connection as a DTE (the default Junos configuration) by removing the **dce** statement from the configuration and the other as a DCE.

Related Documentation

- dce on page 514

 Copyright © 2018, Juniper Networks, Inc.

Receive and Transmit Leaky Bucket Properties Overview

Congestion control is particularly difficult in high-speed networks with high volumes of traffic. When congestion occurs in such a network, it is usually too late to react. You can avoid congestion by regulating the flow of packets into your network. Smoother flows prevent bursts of packets from arriving at (or being transmitted from) the same interface and causing congestion.

For all interface types except ATM, Fast Ethernet, Gigabit Ethernet, and channelized IQ and IQE, you can configure leaky bucket properties, which allow you to limit the amount of traffic received on and transmitted by a particular interface. You effectively specify what percentage of the interface's total capacity can be used to receive or transmit packets. You might want to set leaky bucket properties to limit the traffic flow from a link that is known to transmit high volumes of traffic.

NOTE: Instead of configuring leaky bucket properties, you can limit traffic flow by configuring policers. Policers work on all interfaces. For more information, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*.

The leaky bucket is used at the host-network interface to allow packets into the network at a constant rate. Packets might be generated in a bursty manner, but after they pass through the leaky bucket, they enter the network evenly spaced. In some cases, you might want to allow short bursts of packets to enter the network without smoothing them out. By controlling the number of packets that can accumulate in the bucket, the **threshold** property controls burstiness. The maximum number of packets entering the network in**t** time units is **threshold + rate * t**.

By default, leaky buckets are disabled and the interface can receive and transmit packets at the maximum line rate.

For each DS3 channel on a channelized OC12 interface, you can configure unique receive and transmit buckets.

NOTE: HDLC payload scrambling conflicts with traffic shaping configured using leaky bucket properties. If you configure leaky bucket properties, you must disable payload scrambling, because the Junos OS rejects configurations that have both features enabled. For more information, see *Configuring SONET/SDH HDLC Payload Scrambling for Link Stability*.

Related Documentation

- Configuring Receive and Transmit Leaky Bucket Properties to Reduce Network Congestion on page 156
- *SONET/SDH Interfaces Overview*
- receive-bucket on page 921

- transmit-bucket on page 1040

Configuring Receive and Transmit Leaky Bucket Properties to Reduce Network Congestion

You can configure leaky bucket properties which allow you to limit the amount of traffic received on and transmitted by a particular interface. You can specify what percentage of the interface’s total capacity can be used to receive or transmit packets. You might want to set leaky bucket properties to limit the traffic flow from a link that is known to transmit high volumes of traffic.

To configure leaky bucket properties:

1. In configuration mode, go to the **[edit interfaces *interface-name*]** hierarchy level.

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

2. Configure the **receive-bucket** statement.

   [edit interfaces *interface-name*]
   user@host# **set receive-bucket**

3. Configure the **overflow** option, the **threshold** option, and the **rate** option for the receive leaky bucket, which specifies what percentage of the interface’s total capacity can be used to receive packets.

   [edit interfaces *interface-name* receive-bucket]
   user@host# **set overflow (discard | tag)**
   user@host# **set threshold *bytes***
   user@host# **set rate *percentage***

4. Configure the **transmit-bucket** statement.

   [edit interfaces *interface-name*]
   user@host# **set transmit-bucket**

5. Configure the **overflow** option, the **threshold** option, and the **rate** option for the transmit leaky bucket, which specifies what percentage of the interface’s total capacity can be used to transmit packets.

   [edit interfaces *interface-name* transmit-bucket]
   user@host# **set overflow (discard | tag)**
   user@host# **set threshold *bytes***
   user@host# **set rate *percentage***

Related Documentation

- Receive and Transmit Leaky Bucket Properties Overview on page 155
- *SONET/SDH Interfaces Overview*
- receive-bucket on page 921

 Copyright © 2018, Juniper Networks, Inc.

- transmit-bucket on page 1040

Understanding Unidirectional Traffic Flow on Physical Interfaces

By default, physical interfaces are bidirectional; that is, they both transmit and receive traffic. You can configure unidirectional link mode on a 10-Gigabit Ethernet interface that creates two new physical interfaces that are unidirectional. The new transmit-only and receive-only interfaces operate independently, but both are subordinate to the original parent interface.

The unidirectional interfaces enable the configuration of a unidirectional link topology. Unidirectional links are useful for applications such as broadband video services where almost all traffic flow is in one direction, from the provider to the user. Unidirectional link mode conserves bandwidth by enabling it to be differentially dedicated to transmit and receive interfaces. In addition, unidirectional link mode conserves ports for such applications because the transmit-only and receive-only interfaces act independently. Each can be connected to different routers, for example, reducing the total number of ports required.

NOTE: Unidirectional link mode is currently supported on only the following hardware:

- 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

The transmit-only interface is always operationally up. The operational status of the receive-only interface depends only on local faults; it is independent of remote faults and of the status of the transmit-only interface.

On the parent interface, you can configure attributes common to both interfaces, such as clocking, framing, gigether-options, and sonet-options. On each of the unidirectional interfaces, you can configure encapsulation, MAC address, MTU size, and logical interfaces.

Unidirectional interfaces support IP and IPv6. Packet forwarding takes place by means of static routes and static ARP entries. which you can configure independently on both unidirectional interfaces.

Only transmit statistics are reported on the transmit-only interface (and shown as zero on the receive-only interface). Only receive statistics are reported on the receive-only interface (and shown as zero on the transmit-only interface). Both transmit and receive statistics are reported on the parent interface.

Related Documentation

- unidirectional on page 1051
- Enabling Unidirectional Traffic Flow on Physical Interfaces on page 158

Enabling Unidirectional Traffic Flow on Physical Interfaces

By default, physical interfaces are bidirectional; that is, they both transmit and receive traffic. You can configure unidirectional link mode on a 10-Gigabit Ethernet interface that creates two new physical interfaces that are unidirectional. The new transmit-only and receive-only interfaces operate independently, but both are subordinate to the original parent interface.

To enable unidirectional link mode on a physical interface, perform the following steps:

1. In configuration mode, go to the **[edit interfaces *interface-name*]** hierarchy level:

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

2. Configure the **unidirectional** option to create two new, unidirectional (transmit-only and receive-only) physical interfaces subordinate to the original parent interface.

   ```
   [edit interfaces interface-name]
   user@host# set unidirectional
   ```

NOTE: Unidirectional link mode is currently supported on only the following hardware:

- 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

Related Documentation

- unidirectional on page 1051
- Understanding Unidirectional Traffic Flow on Physical Interfaces on page 157

 Copyright © 2018, Juniper Networks, Inc.

Physical Interface Damping Overview

Physical interface damping limits the advertisement of the up and down transitions (flapping) on an interface. Each time a transition occurs, the interface state is changed, which generates an advertisement to the upper-level routing protocols. Damping helps reduce the number of these advertisements.

From the viewpoint of network deployment, physical interface flaps fall into the following categories:

- Nearly instantaneous multiple flaps of short duration (milliseconds).
- Periodic flaps of long duration (seconds).

Figure 6 on page 159 is used to describe these types of interface flaps and the damping configuration that you can use in each case.

Figure 6: Two Router Interfaces Connected Through Transport Equipment

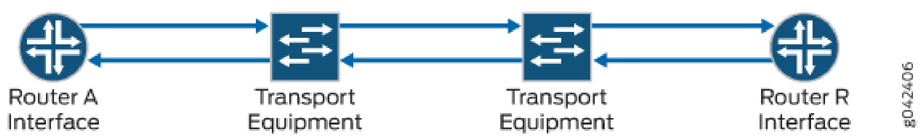

NOTE: We recommend that you use similar damping configurations on both ends of the physical interface. Configuring damping on one end and not having interface damping on the other end can result in undesired behavior.

The following sections describe the types of interface damping depending upon the transition time length.

- Damping Overview for Shorter Physical Interface Transitions on page 159
- Damping Overview for Longer Physical Interface Transitions on page 160

Damping Overview for Shorter Physical Interface Transitions

Figure 6 on page 159 shows two routers with two transport devices between them. If a redundant link between the two transport devices fails, link switching is performed. Link switching takes a number of milliseconds. As shown in Figure 7 on page 160, during switching, both router interfaces might encounter multiple flaps with an up-and-down duration of several milliseconds. These multiple flaps, if advertised to the upper-level routing protocols, might result in undesired route updates. This is why you might want to damp these interface flaps.

NOTE: Damping is suitable only with routing protocols.

Copyright © 2018, Juniper Networks, Inc.

For shorter physical interface transitions, you configure interface damping with the **hold-time** statement on the interface. The hold timer enables interface damping by not advertising interface transitions until the hold timer duration has passed. When a hold-down timer is configured and the interface goes from up to down, the down hold-time timer is triggered. Every interface transition that occurs during the hold-time is ignored. When the timer expires and the interface state is still *down*, then the router begins to advertise the interface as being down. Similarly, when a hold-up timer is configured and an interface goes from down to up, the up hold-time timer is triggered. Every interface transition that occurs during the hold-time is ignored. When the timer expires and the interface state is still *up*, then the router begins to advertise the interface as being up.

Figure 7: Multiple Flaps of Short Duration (Milliseconds)

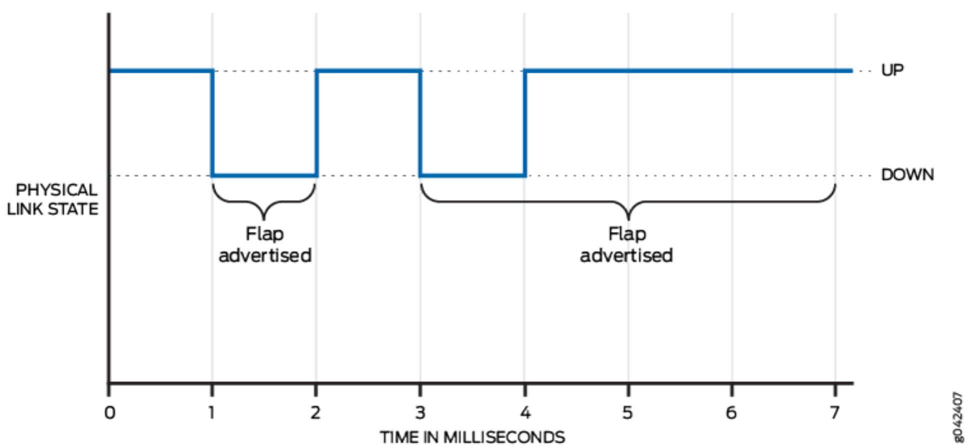

Damping Overview for Longer Physical Interface Transitions

When the link between a router interface and the transport devices is not stable, this can lead to periodic flapping, as shown in Figure 8 on page 161. Flaps occur in the order of seconds or more, with an up-and-down flap duration in the order of a second or more. In this case, using the hold timer feature might not produce optimal results as it cannot suppress the relatively longer and repeated interface flaps. Increasing the hold time duration to seconds still allows the system to send route updates on the flapping interface, so fails to suppress periodically flapping interfaces on the system.

 Copyright © 2018, Juniper Networks, Inc.

Figure 8: Periodic Flaps of Long Duration (Seconds)

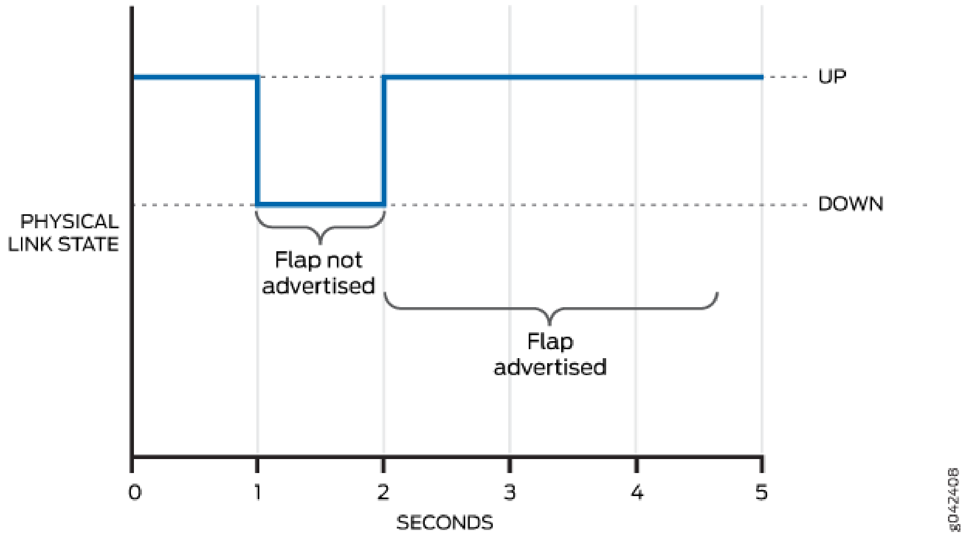

For longer periodic interface flaps, you configure interface damping with the **damping** statement on the interface. This damping method uses an exponential back-off algorithm to suppress interface up-and-down event reporting to the upper-level protocols. Every time an interface goes down, a penalty is added to the interface penalty counter. If at some point the accumulated penalty exceeds the suppress level, the interface is placed in the suppress state, and further interface link up and down events are not reported to the upper-level protocols.

NOTE:

- Only PTX Series routers, T Series routers, MX960 routers, MX480 routers, MX240 routers, MX80 routers, and M10i routers support interface damping for longer periodic interface flaps on all the line cards.
- Penalty added on every interface flap is 1000.
- The system does not indicate whether an interface is down because of suppression or that is the actual state of the physical interface. Because of this, SNMP link traps and Operation, Administration, and Maintenance (OAM) protocols cannot differentiate the damped version of the link state from the real version. Therefore, the traps and protocols might not work as expected.
- You can verify suppression by viewing the information in the **Damping** field of the **show interface extensive** command output.

At all times, the interface penalty counter follows an exponential decay process. Figure 9 on page 163 and Figure 10 on page 164 show the decay process as it applies to recovery when the physical level link is down or up. As soon as the accumulated penalty reaches the lower boundary of the reuse level, the interface is marked as unsuppressed,

Copyright © 2018, Juniper Networks, Inc.

and further changes in the interface link state are again reported to the upper-level protocols. You use the **max-suppress** option to configure the maximum time for restricting the accumulation of the penalty beyond the value of the maximum penalty. The value of the maximum penalty is calculated by the software. The maximum penalty corresponds to the time it would take max-suppress to decay and reach the reuse level. The penalty continues to decay after crossing the reuse level.

Figure 9 on page 163 and Figure 10 on page 164 show the accumulated penalty, and the decay over time as a curve. Whenever the penalty is below the reuse level and the physical level link changes state, state changes are advertised to the system and cause SNMP state changes.

Figure 9 on page 163 shows the penalty dropping below the reuse level when the physical link is down. The system is notified of a state change only after the physical level link transitions to up.

 Copyright © 2018, Juniper Networks, Inc.

Figure 9: Physical-Level Link Is Down When the Penalty Falls Below the Reuse Level

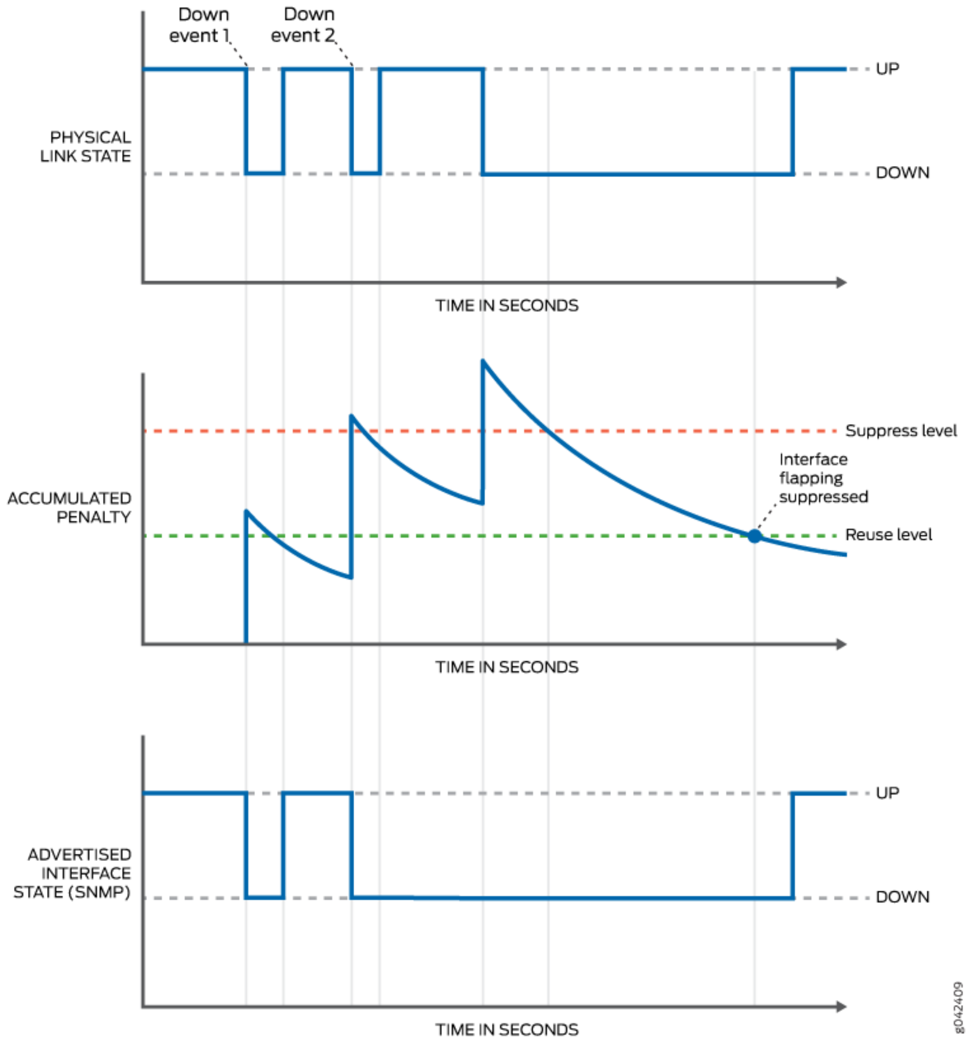

Figure 10 on page 164 shows the penalty dropping below the reuse level when the physical link is up. The system is notified of a state change immediately.

Figure 10: Physical-Level Link Is Up When the Penalty Falls Below the Reuse Level

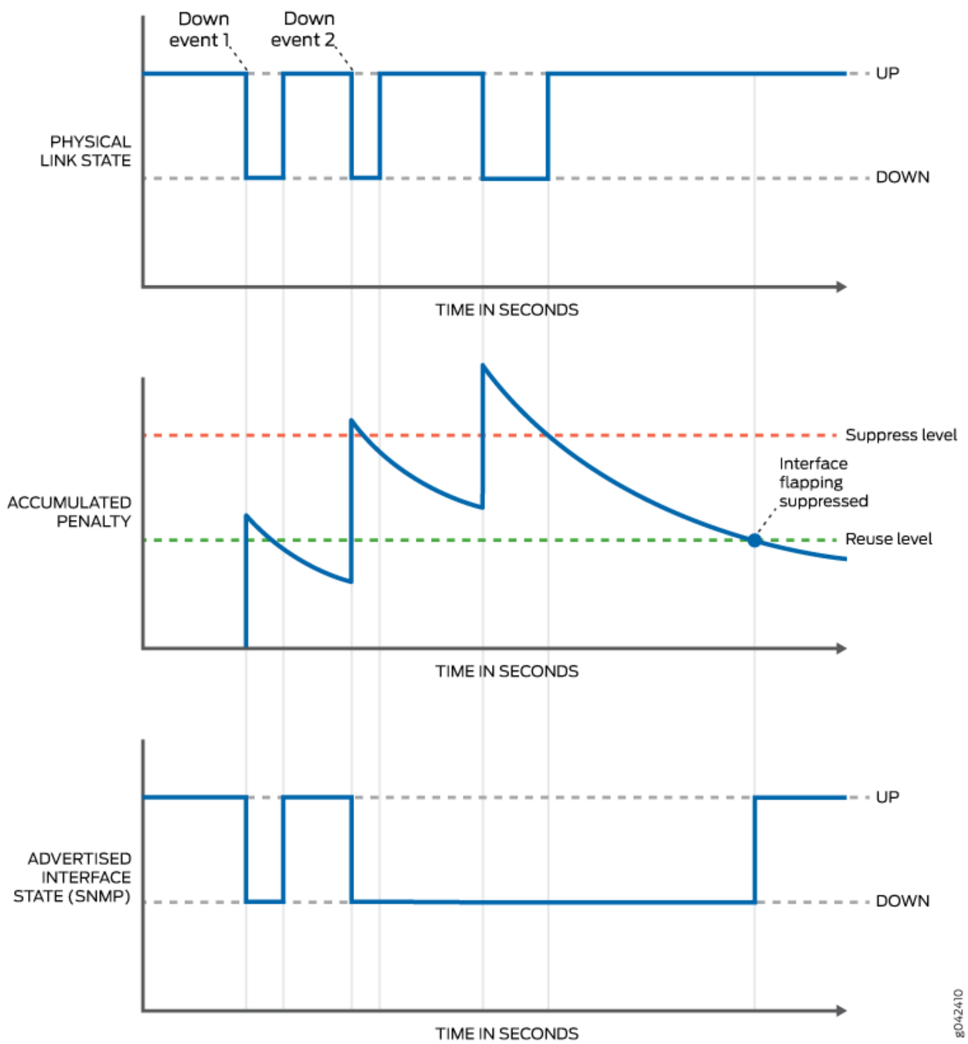

Related Documentation

- Damping Shorter Physical Interface Transitions on page 165
- Damping Longer Physical Interface Transitions on page 166
- *Understanding Damping Parameters*
- **damping on page 509**
- **hold-time on page 644**

Copyright © 2018, Juniper Networks, Inc.

Damping Shorter Physical Interface Transitions

By default, when an interface changes from being up to being down, or from down to up, this transition is advertised immediately to the hardware and Junos OS. In some situations—for example, when an interface is connected to an add/drop multiplexer (ADM) or wavelength-division multiplexer (WDM), or to protect against SONET/SDH framer holes—you might want to damp interface transitions. This means not advertising the interface's transition until a certain period of time has passed, called the *hold-time*. When you have damped interface transitions and the interface goes from up to down, the down hold-time timer is triggered. Every interface transition that occurs during the hold-time is ignored. When the timer expires and the interface state is still *down*, then the router begins to advertise the interface as being down. Similarly, when an interface goes from down to up, the up hold-time timer is triggered. Every interface transition that occurs during the hold-time is ignored. When the timer expires and the interface state is still *up*, then the router begins to advertise the interface as being up. For information about physical interface damping, see "Physical Interface Damping Overview" on page 159.

This task applies to damping shorter physical interface transitions in milliseconds. To damp longer physical interface transitions in seconds, see "Damping Longer Physical Interface Transitions" on page 166.

To configure damping of shorter physical interface transitions:

1. Select the interface to damp, where the interface name is ***interface-type-fpc/pic/port***:

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

2. Configure the hold-time for link up and link down.

   ```
   [edit interfaces interface-name]
   user@host# set hold-time up milliseconds down milliseconds
   ```

The hold time can be a value from 0 through 4,294,967,295 milliseconds. The default value is 0, which means that interface transitions are not damped. Junos OS advertises the transition within 100 milliseconds of the time value you specify.

For most Ethernet interfaces, hold timers are implemented using a one-second polling algorithm. For 1-port, 2-port, and 4-port Gigabit Ethernet interfaces with small form-factor pluggable transceivers (SFPs), hold timers are interrupt-driven.

NOTE: **The hold-time option is not available for controller interfaces.**

Related Documentation

- Physical Interface Damping Overview on page 159
- Damping Longer Physical Interface Transitions on page 166

- *SONET/SDH Defect Hold Times for Damping Interface Transitions Overview*

- *Configuring SONET/SDH Defect Triggers*
- hold-time on page 644

Damping Longer Physical Interface Transitions

Physical interface damping limits the advertisement of the up and down transitions (flapping) on an interface. An unstable link between a router Interface and the transport devices can lead to periodic flapping. Longer flaps occur with a period of about five seconds or more, with an up-and-down duration of one second. For these longer periodic interface flaps, you configure interface damping with the **damping** statement on the interface. This damping method uses an exponential back-off algorithm to suppress interface up and down event reporting to the upper-level protocols. Every time an interface goes down, a penalty is added to the interface penalty counter. If at some point the accumulated penalty exceeds the suppress level **max-suppress**, the interface is placed in the suppress state, and further interface state up and down transitions are not reported to the upper-level protocols.

NOTE:

- Only PTX Series routers, T Series routers, MX2010 routers, MX2020 routers, MX960 routers, MX480 routers, MX240 routers, MX80 routers, and M10i routers support interface damping for longer periodic interface flaps.
- The system does not indicate whether an interface is down because of suppression or that is the actual state of the physical interface. Because of this, SNMP link traps and Operation, Administration, and Maintenance (OAM) protocols cannot differentiate the damped version of the link state from the real version. Therefore, the traps and protocols might not work as expected.
- You can verify suppression by viewing the information in the **Damping** field of the **show interface extensive** command output.

You can view the damping parameters with the **show interfaces extensive** command.

To configure damping of longer physical interface transitions:

1. Select the interface to damp, where the interface name is ***interface-type-fpc/pic/port*** or an interface range:

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

2. Enable longer interface transition damping on a physical interface:

   [edit interfaces *interface-name* damping]
   user@host# **set enable**

Copyright © 2018, Juniper Networks, Inc.

3. (Optional) Set the maximum time in seconds that an interface can be suppressed no matter how unstable the interface has been.

   NOTE: Configure **max-suppress** to a value that is greater than the value of **half-life**; otherwise, the configuration is rejected.

   [edit interfaces *interface-name* damping]
   user@host# **set max-suppress *maximum-seconds***

4. (Optional) Set the decay half-life in seconds, which is the interval after which the accumulated interface penalty counter is reduced by half if the interface remains stable.

   NOTE: Configure **max-suppress** to a value that is greater than the value of **half-life**; otherwise, the configuration is rejected.

   [edit interfaces *interface-name* damping]
   user@host# **set half-life *seconds***

5. (Optional) Set the reuse threshold (no units). When the accumulated interface penalty counter falls below this value, the interface is no longer suppressed.

   [edit interfaces *interface-name* damping]
   user@host# **set reuse *number***

6. (Optional) Set the suppression threshold (no units). When the accumulated interface penalty counter exceeds this value, the interface is suppressed.

   [edit interfaces *interface-name* damping]
   user@host# **set suppress *number***

Related Documentation

- Physical Interface Damping Overview on page 159
- **show interfaces extensive on page 1555**
- Damping Shorter Physical Interface Transitions on page 165
- **damping on page 509**

Example: Configuring Physical Interface Damping

This example shows how to configure damping for a physical interface on a PTX Series Packet Transport Router.

- Requirements on page 168
- Overview on page 168

Requirements

This example uses the following hardware and software components:

- One PTX Series Packet Transport Router
- One or more routers that provide input packets and receive output packets
- Junos OS Release 14.1 or later

Overview

Physical interface damping provides a smoothing of the up and down transitions (flapping) on an interface. Each time a transition occurs, the interface state is changed, which generates an advertisement to the upper-level routing protocols. Damping helps reduce the number of these advertisements.

From the viewpoint of network deployment, physical interface flaps fall into these categories:

- Nearly instantaneous multiple flaps of short duration (milliseconds). For shorter physical interface transitions, you configure interface damping with the **hold-time** statement on the interface. The hold timer enables interface damping by not advertising interface transitions until the hold timer duration has passed. When a hold-down timer is configured and the interface goes from up to down, the interface is not advertised to the rest of the system as being down until it has remained down for the hold-down timer period. Similarly, when a hold-up timer is configured and an interface goes from down to up, it is not advertised as being up until it has remained up for the hold-up timer period.
- Periodic flaps of long duration (seconds). For longer periodic interface flaps, you configure interface damping with the **damping** statement on the interface. This damping method uses an exponential back-off algorithm to suppress interface up and down event reporting to the upper-level protocols. Every time an interface goes down, a penalty is added to the interface penalty counter. If at some point the accumulated penalty exceeds the suppress level, the interface is placed in the suppress state, and further interface state up transitions are not reported to the upper-level protocols.

Configuration

CLI Quick Configuration

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

```
set interfaces xe-6/0/0 damping half-life 11 max-suppress 2222 reuse 3333 suppress
  4444
```

 Copyright © 2018, Juniper Networks, Inc.

Step-by-Step Procedure

To configure damping on the PTX Series Packet Transport Router:

1. Enable damping on the interface, set the half-life interval, maximum suppression, reuse, and suppress values:

   ```
   [edit interface]
   user@router# set xe-6/0/0 damping half-life 11 max-suppress 2222 reuse 3333
     suppress 4444
   ```

2. Commit configuration:

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

Results

From configuration mode, confirm your configuration by entering the **show interfaces** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@router# show interfaces
xe-6 {
  damping {
    half-life 11;
    max-suppress 2222;
    reuse 3333;
    suppress 4444;
  }
```

Verification

To confirm that the configuration is working properly, perform this task:

- Verifying Interface Damping on xe6 on page 169

Verifying Interface Damping on xe6

Purpose

Verify that damping is enabled on the interface and that the damping parameter values are correctly set.

Action From operational mode, run the **show interfaces extensive** command.

```
user@router# run show interfaces xe-6/0/0 extensive
Physical interface: xe-6/0/0, Enabled, Physical link is Up
  Interface index: 158, SNMP ifIndex: 535, Generation: 161
  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
  Damping        : half-life: 11 sec, max-suppress: 2222 sec, reuse: 3333,
suppress: 4444, state: unsuppressed
```

Meaning Damping is enabled and configured successfully on the xe-6 interface.

Related Documentation

- Physical Interface Damping Overview on page 159
- damping on page 509

Configuring Multiservice Physical Interface Properties

The adaptive services (AS), collector, monitoring services, and monitoring services II interfaces are multiservice interfaces specifically designed to enable IP services. To configure multiservice physical interface properties on the collector, monitoring services, and AS interfaces, include the **multiservice-options** statement:

```
multiservice-options {
  (core-dump | no-core-dump);
  (syslog | no-syslog);
  flow-control-options {
    down-on-flow-control;
    dump-on-flow-control;
    reset-on-flow-control;
  }
}
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces cp-*fpc/pic/port*]**
- **[edit interfaces mo-*fpc/pic/port*]**
- **[edit interfaces sp-*fpc/pic/port*]**

For more information about the services interfaces, see the *Junos OS Services Interfaces Library for Routing Devices*.

 Copyright © 2018, Juniper Networks, Inc.

Enabling or Disabling SNMP Notifications on Physical Interfaces

By default, Simple Network Management Protocol (SNMP) notifications are sent when the state of an interface or a connection changes. You can enable or disable these notification based on you requirements.

To explicitly enable sending SNMP notifications on the physical interface, perform the following steps:

1. In configuration mode, go to the **[edit interfaces *interface-name*]** hierarchy level:

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

2. Configure the **traps** option to enable sending of Simple Network Management Protocol (SNMP) notifications when the state of the connection changes.

   ```
   [edit interfaces interface-name]
   user@host# set traps
   ```

To disable sending SNMP notifications on the physical interface, perform the following steps:

1. In configuration mode, go to the **[edit interfaces *interface-name*]** hierarchy level:

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

2. Configure the **no-traps** option to disable sending of Simple Network Management Protocol (SNMP) notifications when the state of the connection changes.

   ```
   [edit interfaces interface-name]
   user@host# set no-traps
   ```

Related Documentation

- traps on page 1044

Configuring Accounting for the Physical Interface

Accounting Profiles Overview

Juniper Networks routers and switches can collect various kinds of data about traffic passing through the router and switch. You can set up one or more *accounting profiles* that specify some common characteristics of this data, including the following:

- The fields used in the accounting records

Copyright © 2018, Juniper Networks, Inc.

- The number of files that the router or switch retains before discarding, and the number of bytes per file
- The polling period that the system uses to record the data

You configure the profiles and define a unique name for each profile using statements at the **[edit accounting-options]** hierarchy level. There are two types of accounting profiles: interface profiles and filter profiles. You configure interface profiles by including the **interface-profile** statement at the **[edit accounting-options]** hierarchy level. You configure filter profiles by including the **filter-profile** statement at the **[edit accounting-options]** hierarchy level. For more information, see the *Network Management and Monitoring Guide*.

You apply filter profiles by including the **accounting-profile** statement at the **[edit firewall filter *filter-name*]** and **[edit firewall family *family* filter *filter-name*]** hierarchy levels. For more information, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*.

Configuring Accounting for the Physical Interface

Before you begin

You must configure a profile to collect error and statistic information for input and output packets on a particular physical interface. An accounting profile specifies what statistics should be collected and written to a log file. For more information on how to configure an accounting-data log file, see the *Configuring Accounting-Data Log Files*.

An interface profile specifies the information collected and written to a log file. You can configure a profile to collect error and statistic information for input and output packets on a particular physical interface.

1. To configure which statistics should be collected for an interface, include the **fields** statement at the **[edit accounting-options interface-profile *profile-name*]** hierarchy level.

   [edit accounting-options interface-profile *profile-name*]
   user@host# **set fields*field-name***

2. Each accounting profile logs its statistics to a file in the **/var/log** directory. To configure which file to use, include the **file** statement at the **[edit accounting-options interface-profile *profile-name*]** hierarchy level.

   [edit accounting-options interface-profile *profile-name*]
   user@host# **set file *filename***

   NOTE: You must specify a **file** statement for the interface profile that has already been configured at the **[edit accounting-options]** hierarchy level. For more information, see the Configuring Accounting-Data Log Files

3. Each interface with an accounting profile enabled has statistics collected once per interval time specified for the accounting profile. Statistics collection time is scheduled

 Copyright © 2018, Juniper Networks, Inc.

evenly over the configured interval. To configure the interval, include the interval statement at the [**edit accounting-options interface-profile *profile-name***] hierarchy level.

```
[edit accounting-options interface-profile profile-name]
user@host# set interval minutes
```

NOTE: The minimum interval allowed is 1 minute. Configuring a low interval in an accounting profile for a large number of interfaces might cause serious performance degradation.

4. To configure the interfaces on which the accounting needs to be performed, apply the interface profile to a physical interface by including the **accounting-profile** statement at the **[edit interfaces interface-name]** hierarchy level.

```
[edit interfaces]
user@host# set interface-name accounting-profile profile-name
```

See Also

- *Accounting Options Overview*
- *Configuring Accounting-Data Log Files*

Displaying Accounting Profile for the Physical Interface

Purpose To display the configured accounting profile a particular physical interface at the **[edit accounting-options interface-profile *profile-name*]** hierarchy level:

- interface-name—ge-1/0/1
- Interface profile —**if_profile**
- File name—**if_stats**
- Interval—15 minutes

Action

- Run the **show** command at the **[edit edit interfaces ge-1/0/1]** hierarchy level.

```
[edit interfaces ge-1/0/1]
accounting-profile if_profile;
```

- Run the **show** command at the **[edit accounting-options]** hierarchy level.

```
interface-profile if_profile {
  interval 15;
  file if_stats {
    fields {
      input-bytes;
      output-bytes;
      input-packets;
      output-packets;
      input-errors;
      output-errors;
```

Copyright © 2018, Juniper Networks, Inc.

```
        }
      }
    }
```

Meaning The configured accounting and its associated set options are displayed as expected.

Disabling a Physical Interface

Disabling a Physical Interface

You can disable a physical interface, marking it as being down, without removing the interface configuration statements from the configuration.

CAUTION: Dynamic subscribers and logical interfaces use physical interfaces for connection to the network. The Junos OS allows you to set the interface to disable and commit the change while dynamic subscribers and logical interfaces are still active. This action results in the loss of all subscriber connections on the interface. Use care when disabling interfaces.

To disable a physical interface:

1. In configuration mode, go to **[edit interfaces *interface-name*]** hierarchy level.

   [edit]
   user@host# **edit interfaces ge-*fpc/pic/port***

2. Include the **disable** statement.

   [edit interfaces at-*fpc/pic/port*]
   user@host# **set disable**

NOTE: On the router, when you use the **disable** statement at the **edit interfaces** hierarchy level, depending on the PIC type, the interface might or might not turn off the laser. Older PIC transceivers do not support turning off the laser, but newer Gigabit Ethernet PICs with SFP and XFP transceivers do support it and the laser will be turned off when the interface is disabled.

WARNING: Do not stare into the laser beam or view it directly with optical instruments even if the interface has been disabled.

 Copyright © 2018, Juniper Networks, Inc.

Example: Disabling a Physical Interface

Sample interface configuration:

```
[edit interfaces]
user@host# show
ge-0/3/2 {
  unit 0 {
    description CE2-to-PE1;
    family inet {
      address 20.1.1.6/24;
    }
  }
}
```

Disabling the interface:

```
[edit interfaces ge-0/3/2]
user@host# set disable
```

Verifying the interface configuration:

```
[edit interfaces ge-0/3/2]
user@host# show
disable; # Interface is marked as disabled.
  unit 0 {
  description CE2-to-PE1;
    family inet {
      address 20.1.1.6/24;
    }
}
```

Effect of Disabling Interfaces on T series PICs

The following table describes the effect of using the **set interfaces disable** ***interface_name*** statement on T series PICs.

Table 47: Effect of set interfaces disable <interface_name> on T series PICs

PIC Model Number	PIC Description	Type of PIC	Behaviour
PF-12XGE-SFPP	10-Gigabit Ethernet LAN/WAN PIC with SFP+ (T4000 Router)	5	Tx laser disabled
PF-24XGE-SFPP	10-Gigabit Ethernet LAN/WAN PIC with Oversubscription and SFP+ (T4000 Router)	5	Tx laser disabled
PF-1CGE-CFP	100-Gigabit Ethernet PIC with CFP (T4000 Router)	5	Tx laser disabled
PD-4XGE-XFP	10-Gigabit Ethernet, 4-port LAN/WAN XFP	4	Tx laser disabled
PD-5-10XGE-SFPP	10-Gigabit LAN/WAN with SFP+	4	Tx laser disabled

Copyright © 2018, Juniper Networks, Inc.

Table 47: Effect of set interfaces disable <interface_name> on T series PICs (continued)

PIC Model Number	PIC Description	Type of PIC	Behaviour
PD-1XLE-CFP	40-Gigabit with CFP	4	Tx laser disabled
PD-1CE-CFP-FPC4	100-Gigabit with CFP	4	Tx laser disabled
PD-TUNNEL	40-Gigabit Tunnel Services	4	NA
PD-4OC192-SON-XFP	OC192/STM64, 4-port XFP	4	Tx laser not disabled
PD-1OC768-SON-SR	OC768c/STM256, 1-port	4	Tx laser not disabled

Related Documentation

- disable on page 537

Copyright © 2018, Juniper Networks, Inc.

CHAPTER 3

Configuring Logical Interface Properties

Logical Interfaces Configuration Properties Overview

For a physical interface device to function, you must configure at least one logical interface on that device. For each logical interface, you must specify the protocol family that the interface supports. You can also configure other logical interface properties. These vary by Physical Interface Card (PIC) and encapsulation type, but include the IP address of the interface, and whether the interface supports multicast traffic, data-link connection

Copyright © 2018, Juniper Networks, Inc.

identifiers (DLCIs), virtual channel identifiers (VCIs) and virtual path identifiers (VPIs), and traffic shaping.

Related Documentation

- Logical Part of an Interface Name on page 39

Logical Interfaces Configuration Statements

To configure logical interface properties, include the following statements:

```
unit logical-unit-number {
  accept-source-mac {
    mac-address mac-address {
      policer {
        input cos-policer-name;
        output cos-policer-name;
      }
    }
  }
  accounting-profile name;
  allow-any-vci;
  atm-scheduler-map (map-name | default);
  backup-options {
    interface interface-name;
  }
  bandwidth rate;
  cell-bundle-size cells;
  clear-dont-fragment-bit;
  compression {
    rtp {
      f-max-period number;
      queues [ queue-numbers ];
      port {
        minimum port-number;
        maximum port-number;
      }
    }
  }
  compression-device interface-name;
  copy-tos-to-outer-ip-header;
  demux-destination family;
  demux-source family;
  demux-options {
    underlying-interface interface-name;
  }
  description text;
  interface {
    l2tp-interface-id name;
    (dedicated | shared);
  }
  dialer-options {
    activation-delay seconds;
    callback;
    callback-wait-period time;
```

 Copyright © 2018, Juniper Networks, Inc.

```
    deactivation-delay seconds;
    dial-string [ dial-string-numbers ];
    idle-timeout seconds;
    incoming-map {
      caller (caller-id| accept-all);
      initial-route-check seconds;
      load-interval seconds;
      load-threshold number;
      pool pool-name;
      redial-delay time;
      watch-list {
        [ routes ];
      }
    }
  }
  disable;
  disable-mlppp-inner-ppp-pfc;
  dlci dlci-identifier;
  drop-timeout milliseconds;
  dynamic-call-admission-control {
    activation-priority priority;
    bearer-bandwidth-limit kilobits-per-second;
  }
  encapsulation type;
  epd-threshold plp1 cells;
  filter filter-name;
  fragment-threshold bytes;
  inner-vlan-id-range start start-id end end-id;
  input-vlan-map {
    inner-tag-protocol-id;
    inner-vlan-id;
    (pop | pop-pop | pop-swap | push | push-push | swap |swap-push | swap-swap);
    tag-protocol-id tpid;
    vlan-id number;
  }
  interleave-fragments;
  inverse-arp;
  link-layer-overhead percent;
  layer2-policer {
    input-policer policer-name;
    input-three-color policer-name;
    output-policer policer-name;
    output-three-color policer-name;
  }
  minimum-links number;
  mrru bytes;
  multicast-dlci dlci-identifier;
  multicast-vci vpi-identifier.vci-identifier;
  multilink-max-classes number;
  multipoint;
  oam-liveness {
    up-count cells;
    down-count cells;
  }
  oam-period (seconds | disable);
  output-vlan-map {
```

```
    inner-tag-protocol-id;
    inner-vlan-id;
    (pop | pop-pop | pop-swap | push | push-push | swap |swap-push | swap-swap);
    tag-protocol-id tpid;
    vlan-id number;
  }
  passive-monitor-mode;
  peer-unit unit-number;
  plp-to-clp;
  point-to-point;
  ppp-options {
    chap {
    access-profile name;
    default-chap-secret name;
    local-name name;
    passive;
  }
  compression {
    acfc;
    pfc;
  }
  dynamic-profile profile-name;
  lcp-restart-timer milliseconds;
  loopback-clear-timer seconds;
  ncp-restart-timer milliseconds;
  pap {
    default-pap-password password;
    local-name name;
    local-password password;
    passive;
  }
  pppoe-options {
    access-concentrator name;
    auto-reconnect seconds;
    (client | server);
    service-name name;
    underlying-interface interface-name;
  }
  proxy-arp;
  service-domain (inside | outside);
  shaping {
    (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate burst
      length);
    queue-length number;
  }
  short-sequence;
  transmit-weight number;
  (traps | no-traps);
  trunk-bandwidth rate;
  trunk-id number;
  tunnel {
    backup-destination address;
    destination address;
    key number;
    routing-instance {
      destination routing-instance-name;
```

 Copyright © 2018, Juniper Networks, Inc.

```
            }
            source source-address;
            ttl number;
        }
        vci vpi-identifier.vci-identifier;
        vci-range start start-vci end end-vci;
        vpi vpi-identifier;
        vlan-id number;
        vlan-id-range number-number;
        vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
        family family {
            [ family-statements ];
        }
    }
```

You can include these statements at the following hierarchy levels:

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

For information about interface-specific logical properties, see Table 48 on page 181.

Logical Interfaces Statements List

Table 48 on page 181 lists statements that you can use to configure logical interfaces.

Table 48: Statements for Logical Interface Properties

Statement	Interface Types	Usage Guidelines
access-profile *name*	ATM2 IQ interfaces	"Configuring the PPP Password Authentication Protocol On a Logical Interface" on page 196
accept-source-mac	Gigabit Ethernet intelligent queuing (IQ) interfaces	*Configuring MAC Address Filtering*
accounting-profile *name*	All	"Configuring Accounting for the Logical Interface" on page 198
allow-any-vci	Asynchronous Transfer Mode (ATM) interfaces	*Configuring ATM Interface Encapsulation*
atm-scheduler-map (*map-name* | default)	ATM2 IQ interfaces	*ATM2 IQ VC Tunnel CoS Components Overview*
backup-destination *address*	Encryption interfaces	*Class of Service Feature Guide for Routing Devices and EX9200 Switches*
bandwidth *rate*	All interfaces, except multilink and aggregated	"Configuring the Interface Bandwidth" on page 189

Table 48: Statements for Logical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
cbr *rate*	ATM interfaces	*Defining the ATM Traffic-Shaping Profile Overview*
cell-bundle-size *cells*	ATM2 IQ interfaces	*Configuring the Layer 2 Circuit Cell-Relay Cell Maximum Overview*
clear-dont-fragment-bit	Adaptive services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
compression	AS PIC or MultiServices PIC link services IQ interfaces (**lsq**) and voice services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
copy-tos-to-outer-ip-header	GRE tunnel interfaces	*Class of Service Feature Guide for Routing Devices and EX9200 Switches*
demux-destination *family*	IP demux interfaces	"Configuring an IP Demultiplexing Interface" on page 316
demux-options *family*	IP demux interfaces	"Configuring an IP Demultiplexing Interface" on page 316
demux-source *family*	IP demux interfaces	"Configuring an IP Demultiplexing Interface" on page 316
description *text*	All	"Adding a Logical Unit Description to the Configuration" on page 188
destination (*address* | *routing-instance-name*)	Encryption generic routing encapsulation (GRE) tunnel, and IP tunnel interfaces	*Junos OS Services Interfaces Library for Routing Devices*
disable	All	"Disabling a Logical Interface" on page 201
disable-mlppp-inner-ppp-pfc	MLPPP interfaces	*Junos OS Services Interfaces Library for Routing Devices*
dlci *dlci-identifier*	Point-to-point interfaces with Frame Relay encapsulation	*Configuring Frame Relay DLCIs*
drop-timeout *milliseconds*	Multilink interfaces	*Junos OS Services Interfaces Library for Routing Devices*
dynamic-profile *profile-name*	1-Gigabit Ethernet and 10-Gigabit Ethernet interfaces configured with PPP over Ethernet on M120 and M320 routers	*Junos Subscriber Access Configuration Guide*
encapsulation *type*	All interfaces, except aggregated SONET/SDH and loopback	*Configuring Interface Encapsulation on Logical Interfaces*

Copyright © 2018, Juniper Networks, Inc.

Table 48: Statements for Logical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
epd-threshold *cells*	ATM2 IQ interfaces	*Configuring the ATM2 IQ EPD Threshold*
f-max-period *number*	AS PIC or MultiServices link services IQ interfaces (**lsq-**) and voice services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
family	All	"Configuring the Protocol Family" on page 206
fragment-threshold *bytes*	Multilink interfaces	*Junos OS Services Interfaces Library for Routing Devices*
inner-tag-protocol-id	Gigabit Ethernet IQ interfaces	*802.1Q VLANs Overview*
inner-vlan-id	Gigabit Ethernet IQ interfaces	*802.1Q VLANs Overview*
inner-vlan-id-range	Gigabit Ethernet, 10-Gigabit Ethernet, and aggregated Ethernet IQ interfaces	"Configuring ATM-to-Ethernet Interworking" on page 282
input	AS PIC or MultiServices link services	*Junos OS Services Interfaces Library for Routing Devices*
input-policer *policer-name*	For 1-Gigabit Ethernet and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces on M Series and T Series routers	*Junos OS Services Interfaces Library for Routing Devices* and *Configuring Gigabit Ethernet Two-Color and Tricolor Policers*
input-three-color *policer-name*	For 1-Gigabit Ethernet and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces on M Series and T Series routers	*Class of Service Feature Guide for Routing Devices and EX9200 Switches* and *Configuring Gigabit Ethernet Two-Color and Tricolor Policers*
input-vlan-map	Gigabit Ethernet IQ interfaces	*Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview*
interleave-fragments	Link services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
inverse-arp	Interfaces with ATM and Frame Relay encapsulation	*Configuring Inverse ATM1 or ATM2 ARP* and *Configuring Inverse Frame Relay ARP*
key *number*	GRE tunnel interfaces on Adaptive Services PICs	*Junos OS Services Interfaces Library for Routing Devices*
layer2-policer	1-Gigabit Ethernet and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces	*Configuring Gigabit Ethernet Two-Color and Tricolor Policers*

Copyright © 2018, Juniper Networks, Inc.

Table 48: Statements for Logical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
lcp-restart-timer	Interfaces with PPP encapsulation	"Configuring the PPP Restart Timers" on page 193
l2tp-interface-id *name*	Adaptive services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
link-layer-overhead *percent*	AS PIC or MultiServices link services IQ interfaces (lsq)	*Junos OS Services Interfaces Library for Routing Devices*
local-name *name*	ATM2 IQ interfaces	"Configuring PPP CHAP Authentication" on page 196 and "Configuring the PPP Password Authentication Protocol On a Logical Interface" on page 196
mac-address *mac-address*	Gigabit Ethernet interfaces and Gigabit Ethernet IQ and IQE interfaces with small form-factor pluggable transceivers (SFPs) (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router)	*Configuring MAC Address Filtering*
minimum-links *number*	Multilink interfaces	*Junos OS Services Interfaces Library for Routing Devices*
mrru *bytes*	Multilink interfaces	*Junos OS Services Interfaces Library for Routing Devices*
multicast-dlci *dlci-identifier*	Point-to-multipoint Frame Relay interfaces	*Configuring a Multicast-Capable Frame Relay Connection*
multicast-vci *vpi-identifiervci-identifier*	Point-to-multipoint ATM1 and ATM2 IQ interfaces	*Configuring the ATM OAM F5 Loopback Cell Threshold*
multilink-max-classes *number*	AS PIC or MultiServices link services IQ interfaces (**lsq-**)	*Junos OS Services Interfaces Library for Routing Devices*
multipoint	All	"Configuring a Multipoint Connection" on page 193
ncp-restart-timer	Interfaces with PPP encapsulation	"Configuring the PPP Restart Timers" on page 193
oam-liveness	ATM1 and ATM2 IQ interfaces	*Configuring the ATM OAM F5 Loopback Cell Threshold*
oam-period (disable | *seconds*)	ATM1 and ATM2 IQ interfaces	*Defining the ATM OAM F5 Loopback Cell Period*

Copyright © 2018, Juniper Networks, Inc.

Table 48: Statements for Logical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
output	All	*Junos OS Services Interfaces Library for Routing Devices*
output-policer *policer-name*	For 1-Gigabit Ethernet and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces on M Series and T Series routers	*Class of Service Feature Guide for Routing Devices and EX9200 Switches* and *Configuring Gigabit Ethernet Two-Color and Tricolor Policers*
output-three-color *policer-name*	For 1-Gigabit Ethernet and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces on M Series and T Series routers	*Class of Service Feature Guide for Routing Devices and EX9200 Switches* and *Configuring Gigabit Ethernet Two-Color and Tricolor Policers*
output-vlan-map	Gigabit Ethernet IQ interfaces	*Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview*
passive (CHAP)	ATM2 IQ interfaces	"Configuring PPP CHAP Authentication" on page 196
passive (PAP)	ATM2 IQ interfaces	"Configuring the PPP Password Authentication Protocol On a Logical Interface" on page 196
passive-monitor-mode	SONET/SDH interfaces	*Enabling Packet Flow Monitoring on SONET/SDH Interfaces*
peer-unit *unit-number*	Logical tunnel interfaces	*Junos OS Services Interfaces Library for Routing Devices*
pfc	Interfaces with PPP, PPP CCC, or PPP TCC encapsulation	"Configuring the PPP Protocol Field Compression" on page 151
plp1 *cells*	ATM2 IQ interfaces	*Configuring the ATM2 IQ EPD Threshold*
plp-to-clp	ATM2 IQ interfaces	*Enabling the PLP Setting to Be Copied to the CLP Bit*
point-to-point	All	"Configuring a Point-to-Point Connection" on page 193
policer	Gigabit Ethernet and Gigabit Ethernet IQ and IQE PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router)	*Configuring MAC Address Filtering*
pop	Gigabit Ethernet IQ interfaces	*Removing a VLAN Tag*
pop-pop	Gigabit Ethernet IQ interfaces	*Removing the Outer and Inner VLAN Tags*

Copyright © 2018, Juniper Networks, Inc.

Table 48: Statements for Logical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
pop-swap	Gigabit Ethernet IQ interfaces	*Removing the Outer VLAN Tag and Rewriting the Inner VLAN Tag*
port	AS PIC or MultiServices or MultiServices link services IQ interfaces (**lsq**) and voice services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
ppp-options	Interfaces with PPP, PPP CCC, or PPP TCC encapsulation	"Configuring PPP CHAP Authentication" on page 196 and "Configuring the PPP Password Authentication Protocol On a Logical Interface" on page 196
proxy-arp	Ethernet interfaces	*Configuring Restricted and Unrestricted Proxy ARP*
push	Gigabit Ethernet IQ interfaces	*Stacking a VLAN Tag*
push-push	Gigabit Ethernet IQ interfaces	*Stacking Two VLAN Tags*
queue-length *number*	ATM1 interfaces	*Configuring the ATM1 Queue Length*
queues [*queue-numbers*]	AS PIC or MultiServices link services IQ interfaces (**lsq**) and voice services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
routing-instance	GRE tunnel and IP tunnel interfaces	*Junos OS Services Interfaces Library for Routing Devices*
rtp	AS PIC or MultiServices link services IQ interfaces (**lsq**) and voice services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
rtvbr peak *rate* sustained *rate* burst *length*	ATM2 interfaces	*Configuring ATM CBR*
service-domain (inside | outside)	Adaptive services interfaces	*Junos OS Services Interfaces Library for Routing Devices*
shaping	ATM1 and ATM2 IQ interfaces	*Defining the ATM Traffic-Shaping Profile Overview*
short-sequence	Multilink interfaces	*Junos OS Services Interfaces Library for Routing Devices*
source *source-address*	Encryption, GRE tunnel, and IP tunnel interfaces	*Junos OS Services Interfaces Library for Routing Devices*
swap	Gigabit Ethernet IQ interfaces	*Configuring Frames with Particular TPIDs to Be Processed as Tagged Frames*

Copyright © 2018, Juniper Networks, Inc.

Table 48: Statements for Logical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
swap-push	Gigabit Ethernet IQ interfaces	*Rewriting a VLAN Tag and Adding a New Tag*
swap-swap	Gigabit Ethernet IQ interfaces	*Rewriting the Inner and Outer VLAN Tags*
tag-protocol-id *tpid*	Gigabit Ethernet and Gigabit Ethernet IQ and IQE PICs with SFPs (except the 10-port Gigabit Ethernet PIC, Aggregated Ethernet with Gigabit Ethernet IQ interfaces, and the built-in Gigabit Ethernet port on the M7i router)	*Rewriting the VLAN Tag on Tagged Frames*
transmit-weight *number*	ATM2 IQ interfaces	*Configuring the ATM2 IQ Transmission Weight*
(traps | no-traps)	All	"Enabling or Disabling SNMP Notifications on Logical Interfaces" on page 200
trunk-bandwidth *rate*	ATM2 IQ interfaces	*Configuring Layer 2 Circuit Trunk Mode Scheduling Overview*
trunk-id *number*	ATM2 IQ interfaces	*Configuring Layer 2 Circuit Transport Mode*
ttl *number*	GRE tunnel and IP tunnel interfaces	*Junos OS Services Interfaces Library for Routing Devices*
tunnel	Encryption, GRE tunnel, and IP tunnel interfaces	*Junos OS Services Interfaces Library for Routing Devices*
underlying-interface	IP demux interfaces	"Configuring an IP Demultiplexing Interface" on page 316
vbr peak *rate* sustained *rate* burst *length*	ATM interfaces	*Defining the ATM Traffic-Shaping Profile Overview*
vci *vpi-identifiervci-identifier*	ATM1 and ATM2 IQ point-to-point interfaces	*Configuring a Point-to-Point ATM1 or ATM2 IQ Connection*
vci-range	ATM2 IQ interfaces	"Configuring ATM-to-Ethernet Interworking" on page 282
vpi *vpi-identifier*	ATM1 and ATM2 IQ point-to-point interfaces	*Configuring a Point-to-Point ATM1 or ATM2 IQ Connection*
vlan-id *number*	Fast Ethernet, Gigabit Ethernet, and Gigabit Ethernet IQ interfaces and aggregated Ethernet using Gigabit Ethernet IQ interfaces	*Binding VLAN IDs to Logical Interfaces* and *Rewriting the VLAN Tag on Tagged Frames*

Copyright © 2018, Juniper Networks, Inc.

Table 48: Statements for Logical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
vlan-tags inner *tpidvlan-id* outer *tpidvlan-id*	Gigabit Ethernet IQ interfaces	*Configuring Dual VLAN Tags*

Specifying the Logical Interface Number

Each logical interface must have a logical unit number. The logical unit number corresponds to the logical unit part of the interface name. For more information, see "Interface Naming Overview" on page 32.

Point-to-Point Protocol (PPP), Cisco High-level Data Link Control (HDLC), and Ethernet circuit cross-connect (CCC) encapsulations support only a single logical interface, whose logical unit number must be 0. Frame Relay and ATM encapsulations support multiple logical interfaces, so you can configure one or more logical unit numbers.

You specify the logical unit number by including the **unit** statement:

```
unit logical-unit-number {
  ...
}
```

You can include this statement at the following hierarchy levels:

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

The logical unit number can be in the range 0 through 65,535 for demux and PPPoE static interfaces only. The logical unit number can be in the range 0 through 16,385 for all other static interface types.

Adding a Logical Unit Description to the Configuration

You can include a text description of each logical unit in the configuration file. Any descriptive text you include is displayed in the output of the **show interfaces** commands, and is also exposed in the **ifAlias** Management Information Base (MIB) object. It has no impact on the interface's configuration. To add a text description, include the **description** statement:

```
description text;
```

You can include this statement at the following hierarchy levels:

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

The description can be a single line of text. If the text contains spaces, enclose it in quotation marks.

 Copyright © 2018, Juniper Networks, Inc.

NOTE: You can configure the extended DHCP relay to include the interface description in the option 82 Agent Circuit ID suboption. See "*Using DHCP Relay Agent Option 82 Information*" in the *Junos OS Broadband Subscriber Management and Services Library*.

For information about describing physical interfaces, see "Configuring Interface Description" on page 113.

Configuring the Interface Bandwidth

By default, the Junos OS uses the physical interface's speed for the MIB-II object, **ifSpeed**. You can configure the logical unit to populate the **ifSpeed** variable by configuring a bandwidth value for the logical interface. The **bandwidth** statement sets an informational-only parameter; you cannot adjust the actual bandwidth of an interface with this statement.

NOTE: We recommend that you be careful when setting this value. Any interface bandwidth value that you configure using the **bandwidth** 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:

cost = reference-bandwidth/bandwidth,

where bandwidth is the physical interface speed. However, if you specify a value for bandwidth using the **bandwidth** statement, that value is used to calculate the interface cost, rather than the actual physical interface bandwidth.

To configure the bandwidth value for a logical interface, include the **bandwidth** statement:

```
bandwidth rate;
```

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*]

rate is the peak rate, in bps or 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 using the formula 1 cps = 384 bps. The value can be any positive integer. The **bandwidth** statement is valid for all logical interfaces, except multilink interfaces.

Copyright © 2018, Juniper Networks, Inc.

Configuring Interface Encapsulation on Logical Interfaces

Understanding Interface Encapsulation on Logical Interfaces

You can configure an encapsulation on a logical interface, which is the encapsulation used within certain packet types.

The following restrictions apply to logical interface encapsulation:

- With the atm-nlpid, atm-cisco-nlpid, and atm-vc-mux encapsulations, you can configure the inet family only.
- With the CCC circuit encapsulations, you cannot configure a family on the logical interface.
- A logical interface cannot have frame-relay-ccc encapsulation unless the physical device also has frame-relay-ccc encapsulation.
- A logical interface cannot have frame-relay-tcc encapsulation unless the physical device also has frame-relay-tcc encapsulation. In addition, you must assign this logical interface a DLCI from 512 through 1022 and configure it as point-to-point.
- A logical interface cannot have frame-relay-ether-type or frame-relay-ether-type-tcc encapsulation unless the physical interface has flexible-frame-relay encapsulation and is on an IQ or IQE PIC.
- For frame-relay-ether-type-tcc encapsulation, you must assign this logical interface a DLCI from 512 through 1022.
- For interfaces that carry IP version 6 (IPv6) traffic, you cannot configure ether-over-atm-llc encapsulation.
- When you use ether-over-atm-llc encapsulation, you cannot configure multipoint interfaces.
- A logical interface cannot have vlan-ccc or vlan-vpls encapsulation unless the physical device also has vlan-ccc or vlan-vpls encapsulation, respectively. In addition, you must assign this logical interface a VLAN ID from 512 through 1023; if the VLAN ID is 511 or lower, it is subject to the normal destination filter lookups in addition to source address filtering. For more information, see *Configuring VLAN and Extended VLAN Encapsulation*.
- You can create an ATM cell-relay circuit by configuring an entire ATM physical device or an individual virtual circuit (VC). When you configure an entire device, only cell-relay encapsulation is allowed on the logical interfaces. For more information, see *Configuring an ATM1 Cell-Relay Circuit Overview*.

 Copyright © 2018, Juniper Networks, Inc.

Configuring the Encapsulation on a Logical Interface

Generally, you configure an interface's encapsulation at the **[edit interfaces *interface-name*]** hierarchy level. However, for some encapsulation types, such as Frame Relay, ATM, and Ethernet virtual local area network (VLAN) encapsulations, you can also configure the encapsulation type that is used inside the Frame Relay, ATM, or VLAN circuit itself.

To configure encapsulation on a logical interface:

1. In configuration mode, go to the **[edit interfaces *interface-name* unit *logical-unit-number*]** or **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

   [edit]
   user@host# **set interfaces at-*fpc/pic/port* unit *logical-unit-number***

2. Configure the encapsulation type as described in encapsulation (Logical Interface).

   [edit interfaces at-*fpc/pic/port* unit *logical-unit-number*]
   user@host# **set encapsulation *encapsulation-type***

Displaying the Encapsulation on a Logical Interface

Purpose To display the configured encapsulation and its associated set options on a physical interface when the following are set at the **[edit interfaces *interface-name*]** or **[edit logical-systems *logical-system-name* interfaces *interface-name*]** hierarchy level:

- interface-name—at-1/1/0
- Encapsulation—**atm-ccc-cell-relay**
- Unit—120

Action Run the **show** command at the **[edit interfaces *interface-name*]** hierarchy level.

```
[edit interfaces at-1/1/0]
user@host# show
encapsulation atm-ccc-cell-relay;
unit 120 {
  encapsulation atm-ccc-cell-relay;
}
```

Meaning The configured encapsulation and its associated set options are displayed as expected.

Related Documentation

- encapsulation (Logical Interface) on page 573
- *Configuring VLAN and Extended VLAN Encapsulation*
- *Configuring an ATM1 Cell-Relay Circuit Overview*

Copyright © 2018, Juniper Networks, Inc.

Configuring Interface Encapsulation on PTX Series Packet Transport Routers

This topic describes how to configure interface encapsulation on PTX Series Packet Transport Routers. Use the **flexible-ethernet-services** configuration statement to configure different encapsulation for different logical interfaces under a physical interface. With flexible Ethernet services encapsulation, you can configure each logical interface encapsulation without range restrictions for VLAN IDs.

Supported encapsulations for physical interfaces include:

- **flexible-ethernet-services**
- **ethernet-ccc**
- **ethernet-tcc**

Supported encapsulations for logical interfaces include:

- **ethernet**
- **vlan-ccc**
- **vlan-tcc**

NOTE: PTX Series Packet Transport Routers do not support **extended-vlan-cc** and **extended-vlan-tcc** encapsulation on logical interfaces. Instead, you can configure a tag protocol ID (TPID) value of 0x9100 to achieve the same results.

To configure flexible Ethernet services encapsulation, include the **encapsulation flexible-ethernet-services** statement at the **[edit interfaces et-*fpc/pic/port*]** hierarchy level. For example:

```
interfaces {
  et-fpc/pic/port {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
      vlan-id 1000;
      family inet {
        address 11.0.0.20/24;
      }
    }
    unit 1 {
      encapsulation vlan-ccc;
      vlan-id 1010;
    }
    unit 2 {
      encapsulation vlan-tcc;
      vlan-id 1020;
      family tcc {
        proxy {
```

 Copyright © 2018, Juniper Networks, Inc.

```
                inet-address 11.0.2.160;
              }
              remote {
                inet-address 11.0.2.10;
              }
            }
          }
        }
      }
```

Related Documentation

- Configuring Interface Encapsulation on Physical Interfaces on page 129

Configuring a Point-to-Point Connection

By default, all interfaces are assumed to be point-to-point connections. You must ensure that the maximum transmission unit (MTU) sizes on both sides of the connection are the same.

For all interfaces except aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet, you can explicitly configure an interface to be a point-to-point connection by including the **point-to-point** statement:

```
point-to-point;
```

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*]**

Configuring a Multipoint Connection

By default, all interfaces are assumed to be point-to-point connections. To configure an interface to be a multipoint connection, include the **multipoint** statement:

```
multipoint;
```

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*]**

Configuring the PPP Restart Timers

You can configure a restart timer for the Link Control Protocol (LCP) and Network Control Protocol (NCP) components of a PPP session. You can configure the LCP restart timer on interfaces with PPP, PPP TCC, PPP over Ethernet, PPP over ATM, and PPP over Frame Relay encapsulations. You can configure the NCP restart timer on interfaces with PPP and PPP TCC encapsulations and on multilink PPP bundle interfaces.

Copyright © 2018, Juniper Networks, Inc.

To configure the restart timer for the NCP component of a PPP session, include the **ncp-restart-timer** statement, and specify the number of milliseconds.

To configure the restart timer for the LCP component of a PPP session, include the **lcp-restart-timer** statement, and specify the number of milliseconds:

lcp-restart-timer *milliseconds*;
ncp-restart-timer *milliseconds*;

You can include these statements at the following hierarchy levels:

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

To monitor the configuration, issue the **show interfaces *interface-name*** command. Configured options are displayed in the **PPP parameters** field for the physical interface.

```
user@host> run show interfaces t1-0/0/0:1:1.0 detail
   Logical interface t1-0/0/0:1:1.0 (Index 67) (SNMP ifIndex 40)
   (Generation 156)
     Flags: Hardware-Down Device-Down Point-To-Point SNMP-Traps 0x4000
     Encapsulation: PPP
     PPP parameters:
       LCP restart timer: 2000 msec
       NCP restart timer: 2000 msec
     Protocol inet, MTU: 1500, Generation: 163, Route table: 0
       Flags: Protocol-Down
       Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
         Destination: 1.1.1/24, Local: 1.1.1.2, Broadcast: 1.1.1.255,
```

Configuring the PPP Clear Loop Detected Timer

When a Point-to-Point Protocol (PPP) session detects a loop, the loop detected flag is set. If the flag is not cleared by the protocol after the loopback is cleared, the clear loop detected timer clears the flag after the specified time has elapsed.

To configure the clear loop detected timer for the LCP component of a PPP session, include the **loopback-clear-timer** statement, and specify the number of seconds.

loopback-clear-timer *seconds*;

You can include this statement at the following hierarchy levels:

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

To monitor the configuration, issue the **show interfaces *interface-name* extensive** command.

 Copyright © 2018, Juniper Networks, Inc.

Configuring the LCP Configure-Request Maximum Sent

Link Control Protocol (LCP) Configure-Request is used to establish a link. You can configure the maximum number of LCP Configure-Requests to send. The router stops sending LCP Configure-Requests after the specified maximum number is sent. To configure the LCP Configure-Request maximum, use the **lcp-max-conf-req** statement at the **[edit interfaces *interface-name* unit *number* ppp-options]** hierarchy level. The ***number*** range is from 0 to 65,535; where 0 specifies no limit and the LCP Configure-Request is sent indefinitely. The default is 254.

Related Documentation

- lcp-max-conf-req on page 719

Configuring the NCP Configure-Request Maximum Sent

Network Control Protocol (NCP) Configure-Request is used to establish a link. You can configure the maximum number of NCP Configure-Requests to send. The router stops sending NCP Configure-Requests after the specified maximum number is sent. To configure the NCP Configure-Request maximum, use the **ncp-max-conf-req** statement at the **[edit interfaces *interface-name* unit *number* ppp-options]** hierarchy level. The ***number*** range is from 0 to 65,535; where 0 specifies no limit and NCP Configure-Request is sent indefinitely. The default is 254.

Related Documentation

- ppp-options on page 891

Configuring Dynamic Profiles for PPP

A dynamic profile acts as a template that enables you to create, update, or remove a configuration that includes attributes for client access (for example, interface or protocol) or service (for example, IGMP). Using these profiles you can consolidate all of the common attributes of a client (and eventually a group of clients) and apply the attributes simultaneously.

After they are created, the profiles reside in a profile library on the router. You can then use the **dynamic-profile** statement to attach profiles to interfaces. To assign a dynamic profile to a PPP interface, you can include the **dynamic-profile** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* ppp-options]** hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number ppp-options]
dynamic-profile profile-name;
```

To monitor the configuration, issue the **show interfaces *interface-name*** command.

For information about dynamic profiles, see *Dynamic Profiles Overview* in the *Junos Subscriber Access Configuration Guide*.

For information about creating dynamic profiles, see *Configuring a Basic Dynamic Profile* in the *Junos Subscriber Access Configuration Guide*.

Copyright © 2018, Juniper Networks, Inc.

For information about assigning a dynamic profile to a PPP interface, see *Attaching Dynamic Profiles to Static PPP Subscriber Interfaces* in the *Junos Subscriber Access Configuration Guide*.

NOTE: Dynamic profiles for PPP subscribers are supported only on PPPoE interfaces for this release.

Related Documentation

- *Configuring Dynamic Authentication for PPP Subscribers*

Configuring PPP CHAP Authentication

For interfaces with PPP encapsulation, you can configure interfaces to support the PPP Challenge Handshake Authentication Protocol (CHAP), as defined in RFC 1994, *PPP Challenge Handshake Authentication Protocol (CHAP)*. When you enable CHAP on an interface, the interface can authenticate its peer and can be authenticated by its peer.

For information about configuring CHAP, see “Configuring the PPP Challenge Handshake Authentication Protocol” on page 136.

Configuring the PPP Password Authentication Protocol On a Logical Interface

Understanding PPP Password Authentication Protocol

The Password Authentication Protocol (PAP) provides a simple method for the peer to establish its identity using a two-way handshake. This is done only upon initial link establishment.

After the link is established, an ID and password pair is repeatedly sent by the peer to the authenticator until authentication is acknowledged or the connection is terminated.

For interfaces with PPP encapsulation, you can configure interfaces to support the Password Authentication Protocol (PAP), as defined in RFC 1334, *PAP Authentication Protocols*. If authentication is configured, the PPP link negotiates using CHAP or PAP protocol for authentication during the Link Control Protocol (LCP) negotiation phase. PAP is only performed after the link establishment phase (LCP up) portion of the authentication phase.

During authentication, the PPP link sends a PAP authentication-request packet to the peer with an ID and password. The authentication-request packet is sent every 2 seconds, similar to the CHAP challenge, until a response is received (acknowledgment packet, nonacknowledgment packet). If an acknowledgment packet is received, the PPP link transitions to the next state, the network phase. If a nonacknowledgment packet is received, an LCP terminate request is sent, and the PPP link goes back to the link

 Copyright © 2018, Juniper Networks, Inc.

establishment phase. If no response is received, and an optional retry counter is set to **true**, a new request acknowledgment packet is resent. If the retry counter expires, the PPP link transitions to the LCP negotiate phrase.

You can configure the PPP link with PAP in passive mode. By default, when PAP is enabled on an interface, the interface expects authenticate-request packets from the peer. However, the interface can be configured to send authentication request packets to the peer by configuring PAP to operate in passive mode. In PAP passive mode, the interface sends the authenticate-request packets to the peer only if the interface receives the PAP option from the peer during LCP negotiation—in passive mode, the interface does not authenticate the peer.

Configuring the PPP Password Authentication Protocol On a Logical Interface

When you configure an interface to use PAP, you must assign an access profile to the interface. When an interface receives PAP authentication requests, the access profile in the packet is used to look up the password. If no matching access profile is found for the PAP authentication request that was received by the interface, the optionally configured default PAP password is used.

To configure the PPP password authentication protocol, on each logical interface with PPP encapsulation, perform the following steps.

1. The default PAP password is used when no matching PAP access profile exists, or if the PAP access profile name changes during PPP link negotiation. To configure the default PAP password, include the **pap-password** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* ppp-options pap]** hierarchy level:

   [edit interfaces *interface-name* unit *logical-unt-number* ppp-options pap]
   user@host# **set** default-pap-password ***password***

2. To configure the name the interface uses in PAP request and response packets, include the **local-name** statement at the **[edit interfaces *interface-name* unit *logical-unt-number* ppp-options pap]** hierarchy level:

   [edit interfaces *interface-name* ppp-options pap]
   user@host# **set** local-name ***name***

   NOTE: By default, when PAP is enabled on an interface, the interface uses the router's system hostname as the name sent in PAP request and response packets.

3. You need to configure the password to be used for authentication. To configure the host password for sending PAP requests, include the **local-password** statement at the **[edit interfaces *interface-name* ppp-options pap]** hierarchy level:

   [edit interfaces *interface-name* unit *logical-unt-number* ppp-options pap]
   user@host# **set** local-password ***password***

4. To configure the interface to authenticate with PAP in passive mode, include the **passive** statement at the **[edit interfaces *interface-name* unit *logical-unt-number*ppp-options pap]** hierarchy level:

```
[edit interfaces interface-name unit logical-unt-number ppp-options pap]
user@host# set passive
```

NOTE: By default, when PAP is enabled on an interface, the interface expects authenticate-request packets from the peer. However, the interface can be configured to send authentication request packets to the peer by configuring PAP to operate in passive mode. In PAP passive mode, the interface sends the authenticate-request packets to the peer only if the interface receives the PAP option from the peer during LCP negotiation—in passive mode, the interface does not authenticate the peer.

See Also Configuring the PPP Authentication Protocol on page 142

Configuring Accounting for the Logical Interface

Accounting Profiles Overview

Juniper Networks routers and switches can collect various kinds of data about traffic passing through the router and switch. You can set up one or more *accounting profiles* that specify some common characteristics of this data, including the following:

- The fields used in the accounting records
- The number of files that the router or switch retains before discarding, and the number of bytes per file
- The polling period that the system uses to record the data

You configure the profiles and define a unique name for each profile using statements at the **[edit accounting-options]** hierarchy level. There are two types of accounting profiles: interface profiles and filter profiles. You configure interface profiles by including the **interface-profile** statement at the **[edit accounting-options]** hierarchy level. You configure filter profiles by including the **filter-profile** statement at the **[edit accounting-options]** hierarchy level. For more information, see the *Network Management and Monitoring Guide*.

You apply filter profiles by including the **accounting-profile** statement at the **[edit firewall filter *filter-name*]** and **[edit firewall family *family* filter *filter-name*]** hierarchy levels. For more information, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*.

 Copyright © 2018, Juniper Networks, Inc.

Configuring Accounting for the Logical Interface

Before you begin

You must configure a profile to collect error and statistic information for input and output packets on a particular logical interface. An accounting profile specifies what statistics should be collected and written to a log file. For more information on how to configure an accounting-data log file, see the *Configuring Accounting-Data Log Files*.

An interface profile specifies the information collected and written to a log file. You can configure a profile to collect error and statistic information for input and output packets on a particular logical interface.

1. To configure which statistics should be collected for an interface, include the **fields** statement at the **[edit accounting-options interface-profile *profile-name*]** hierarchy level.

   [edit accounting-options interface-profile *profile-name*]
   user@host# **set fields*field-name***

2. Each accounting profile logs its statistics to a file in the **/var/log** directory. To configure which file to use, include the **file** statement at the **[edit accounting-options interface-profile *profile-name*]** hierarchy level.

   [edit accounting-options interface-profile *profile-name*]
   user@host# **set file *filename***

   NOTE: You must specify a file statement for the interface profile that has already been configured at the **[edit accounting-options]** hierarchy level. For more information, see the Configuring Accounting-Data Log Files

3. Each interface with an accounting profile enabled has statistics collected once per interval time specified for the accounting profile. Statistics collection time is scheduled evenly over the configured interval. To configure the interval, include the interval statement at the [**edit accounting-options interface-profile *profile-name***] hierarchy level.

   [edit accounting-options interface-profile *profile-name*]
   user@host# **set interval *minutes***

   NOTE: The minimum interval allowed is 1 minute. Configuring a low interval in an accounting profile for a large number of interfaces might cause serious performance degradation.

4. To configure the interfaces on which the accounting needs to be performed, apply the interface profile to a logial interface by including the **accounting-profile** statement at the **[edit interfaces interface-name unit *logical-unit-number*]** hierarchy level.

Copyright © 2018, Juniper Networks, Inc.

[edit interfaces]
user@host# set *interface-name* unit *logical-unit-number*accounting-profile *profile-name*

See Also

- *Accounting Options Overview*
- *Configuring Accounting-Data Log Files*

Displaying Accounting Profile for the Logical Interface

Purpose To display the configured accounting profile a particular logical interface at the **[edit accounting-options interface-profile *profile-name*]** hierarchy level:

- interface-name—ge-1/0/1
- Logical unit number—1
- Interface profile —**if_profile**
- File name—**if_stats**
- Interval—15 minutes

Action

- Run the **show** command at the **[edit interfaces ge-1/0/1 unit 1]** hierarchy level.

```
[edit interfaces ge-1/0/1 unit 1]
accounting-profile if_profile;
```

- Run the **show** command at the **[edit accounting-options]** hierarchy level.

```
interface-profile if_profile {
  interval 15;
  file if_stats {
    fields {
      input-bytes;
      output-bytes;
      input-packets;
      output-packets;
      input-errors;
      output-errors;
    }
  }
}
```

Meaning The configured accounting and its associated set options are displayed as expected.

Enabling or Disabling SNMP Notifications on Logical Interfaces

By default, Simple Network Management Protocol (SNMP) notifications are sent when the state of an interface or a connection changes. To explicitly enable these notifications on the logical interface, include the **traps** statement; to disable these notifications on the logical interface, include the **no-traps** statement:

```
(traps | no-traps);
```


Copyright © 2018, Juniper Networks, Inc.

You can include these statements at the following hierarchy levels:

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

Disabling a Logical Interface

You can unconfigure a logical interface, effectively disabling that interface, without removing the logical interface configuration statements from the configuration. To do this, include the **disable** statement:

```
disable;
```

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*****]**

When an interface is disabled, a route (pointing to the reserved target “**REJECT**”) with the IP address of the interface and a 32–bit subnet mask is installed in the routing table. See *Routing Protocols*.

Configuring Logical System Interface Properties

With Junos OS, you can partition a single physical router into multiple logical devices that perform independent routing tasks. Because logical systems perform a subset of the tasks once handled by the physical router, logical systems offer an effective way to maximize the use of a single router.

1. Configure the physical interface that needs to be partitioned into multiple logical systems.

   ```
   [edit]
   user@host# set interfaces interface-name  description description
   ```

2. Create the logical system interface on the logical unit.

   ```
   [edit]
   user@host# set logical-systems name interfaces interface-name  unit
     logical-unit-number description description
   ```

3. Configure the required properties for the logical system.

   ```
   [edit logical-systems name]
   user@host# set interfaces interface-name  unit logical-unit-number family name
     address address
   ```

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- *Examples: Using Logical Systems*

Copyright © 2018, Juniper Networks, Inc.

CHAPTER 4

Configuring Protocol Family and Interface Address Properties

Protocol Family Configuration and Interface Address Statements

For each logical interface, you must configure one or more protocol families. You can also configure interface address properties. To do this, include the following statements:

```
family family {
  accounting {
```

```
  destination-class-usage;
  source-class-usage {
    direction;
  }
}
address address {
  destination address;
}
bundle interface-name;
filter {
  dialer filter-name;
  input filter-name;
  output filter-name;
  group filter-group-number;
}
interface-mode (access | trunk);
ipsec-sa sa-name;
keep-address-and-control;
mtu bytes;
multicast-only;
negotiate-address;
no-redirects:
policer {
  arp policer-template-name;
  input policer-template-name;
  output policer-template-name;
}
primary;
protocols [inet iso mpls];
proxy inet-address address;
receive-options-packets;
receive-ttl-exceeded;
remote (inet-address address | mac-address address);
rpf-check <fail-filter filter-name>;
sampling {
  direction;
}
service {
  input {
    service-set service-set-name <service-filter filter-name>;
    post-service-filter filter-name;
  }
  output {
    service-set service-set-name <service-filter filter-name>;
  }
}
targeted-broadcast {
  forward-and-send-to-re;
  forward-only;
}
(translate-discard-eligible | no-translate-discard-eligible);
(translate-fecn-and-becn | no-translate-fecn-and-becn);
translate-plp-control-word-de;
vlan-id number;
vlan-id-list [number number-number];
```

 Copyright © 2018, Juniper Networks, Inc.

```
    unnumbered-address interface-name destination address destination-profile
      profile-name;
    address address {
      arp ip-address (mac | multicast-mac) mac-address <publish>;
      broadcast address;
      destination address;
      destination-profile name;
      eui-64;
      multipoint-destination address dlci dlci-identifier;
      multipoint-destination address {
        epd-threshold cells;
        inverse-arp;
        oam-liveness {
          up-count cells;
          down-count cells;
        }
        oam-period (disable | seconds);
        shaping {
          (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate
            burst length);
          queue-length number;
        }
        vci vpi-identifier.vci-identifier;
      }
      primary;
      preferred;
      (vrrp-group | vrrp-inet6-group) group-number {
        (accept-data | no-accept-data);
        advertise-interval seconds;
        authentication-type authentication;
        authentication-key key;
        fast-interval milliseconds;
        (preempt | no-preempt) {
          hold-time seconds;
        }
        priority-number number;
        track {
          priority-cost seconds;
          priority-hold-time interface-name {
            interface priority;
            bandwidth-threshold bits-per-second {
              priority;
            }
          }
          route ip-address/mask routing-instance instance-name priority-cost cost;
        }
        virtual-address [ addresses ];
      }
    }
  }
```

You can include these statements at the following hierarchy levels:

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

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

For information about interface-specific protocol and address properties, see *Configuring T1 and NxDS0 Interfaces*.

Related Documentation
- *Configuring T1 and NxDS0 Interfaces*

Configuring the Protocol Family

A protocol family is a group of logical properties within an interface configuration. Protocol families include all the protocols that make up a protocol suite. To use a protocol within a particular suite, you must configure the entire protocol family as a logical property for an interface.

Junos OS protocol families include the following common protocol suites:

- Inet—Supports IP protocol traffic, including OSPF, BGP, and Internet Control Message Protocol (ICMP).
- Inet6—Supports IPv6 protocol traffic, including RIP for IPv6 (RIPng), IS-IS, and BGP.
- ISO—Supports IS-IS traffic.
- MPLS—Supports MPLS.

In addition to the common protocol suites, JUNOS protocol families sometimes use the following protocol suites. For more information see, family.

To configure the logical interface's protocol family, include the **family** statement, specifying the selected family. To configure the protocol family, following are the minimum configuration tasks under the **[edit interfaces interface-name unit logical-unit-number family family]** hierarchy.

Table 49: Protocol Family Configuration Tasks

Task	Find Details Here
Configure MTU	"Configuring the Media MTU" on page 115
Configure the unit and family so that the interface can transmit and receive multicast traffic only	*Restricting Tunnels to Multicast Traffic*
Disable the sending of redirect messages by the router	*Configuring Junos OS to Disable Protocol Redirect Messages on the Router or Switch*
Assign an address to an interface	"Configuring the Interface Address" on page 207

Related Documentation
- family on page 607

Copyright © 2018, Juniper Networks, Inc.

Configuring the Interface Address

You assign an address to an interface by specifying the address when configuring the protocol family. For the **inet** or **inet6** family, configure the interface IP address. For the **iso** family, configure one or more addresses for the loopback interface. For the **ccc**, **ethernet-switching**, **tcc**, **mpls**, **tnp**, and **vpls** families, you never configure an address.

NOTE: The point-to-point (PPP) address is taken from the loopback interface address that has the primary attribute. When the loopback interface is configured as an unnumbered interface, it takes the primary address from the donor interface.

To assign an address to an interface, perform the following steps:

Copyright © 2018, Juniper Networks, Inc.

1. Configure the interface address at the **[edit interfaces *interface-name* unit *logical-unit-number* family *family*]** hierarchy level.

   - To configure an IPv4 address on routers and switches running Junos OS, use the **interface *interface-name* unit *number* family inet address *a.b.c.d/nn*** statement at the **[edit interfaces]** hierarchy level.

     [edit interfaces]
     user@host# **set *interface-name* unit *logical-unit-number* family inet address *a.b.c.d/nn***

     NOTE:
     - Juniper Networks routers and switches support /31 destination prefixes when used in point-to-point Ethernet configurations; however, they are not supported by many other devices, such as hosts, hubs, routers, or switches. You must determine if the peer system also supports /31 destination prefixes before configuration.
     - You can configure the same IPv4 address on multiple physical interfaces. When you assign the same IPv4 address to multiple physical interfaces, the operational behavior of those interfaces differs, depending on whether they are implicitly or explicitly point-to-point .
     - By default, all interfaces are assumed to be point-to-point (PPP) interfaces. For all interfaces except aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet, you can explicitly configure an interface to be a point-to-point connection.
     - If you configure the same IP address on multiple interfaces in the same routing instance, Junos OS uses only the first configuration. The remaining IP address configurations are ignored, leaving some interfaces without an assigned address. Interfaces without an assigned address cannot be used as a donor interface for an unnumbered Ethernet interface.

   - To configure an IPv6 address on routers and switches running Junos OS, use the **interface *interface-name* unit *number* family inet6 address *aaaa:bbbb:...:zzzz/nn*** statement at the **[edit interfaces]** hierarchy level.

     [edit interfaces]
     user@host# **set *interface-name* unit *logical-unit-number* family inet6 address *aaaa:bbbb:...:zzzz/nn***

     NOTE:
     - You represent IP version 6 (IPv6) addresses in hexadecimal notation using a colon-separated list of 16-bit values. The double colon (::) represents all bits set to 0.
     - You must manually configure the router or switch advertisement and advertise the default prefix for autoconfiguration to work on a specific interface.

Copyright © 2018, Juniper Networks, Inc.

2. [Optional] Set the broadcast address on the network or subnet .

   [edit interfaces *interface-name* unit *logical-unit-number* family *family* address *address*],
   user@host# **set broadcast *address***

   NOTE: The broadcast address must have a host portion of either all ones or all zeros. You cannot specify the addresses 0.0.0.0 or 255.255.255.255

3. [Optional] specify the remote address of the connection for the encrypted, PPP-encapsulated, and tunnel interfaces.

   [edit logical-systems *logical-system-name* interfaces *interface-name* unit
   logical-unit-number family *family* address *address*]
   user@host# **set destination *address***

4. [Optional] For interfaces that carry IP version 6 (IPv6) traffic, configure the host to assign iteslf a unique 64-Bit IP Version 6 interface identifier (EUI-64).

   [edit logical-systems *logical-system-name* interfaces *interface-name* unit
   logical-unit-number family *family* address *address*]
   user@host# **set eui-64**

Related Documentation

- Configuring Default, Primary, and Preferred Addresses and Interfaces on page 209

Configuring Default, Primary, and Preferred Addresses and Interfaces

Default, Primary, and Preferred Addresses and Interfaces

The router has a default address and a primary interface, and interfaces have primary and preferred addresses.

The *default address* of the router is used as the source address on unnumbered interfaces. The routing protocol process tries to pick the default address as the router ID, which is used by protocols, including OSPF and internal BGP (IBGP).

The *primary interface* for the router is the interface that packets go out when no interface name is specified and when the destination address does not imply a particular outgoing interface.

An interface's *primary address* is used by default as the local address for broadcast and multicast packets sourced locally and sent out the interface. An interface's *preferred address* is the default local address used for packets sourced by the local router to destinations on the subnet.

Copyright © 2018, Juniper Networks, Inc.

The default address of the router is chosen using the following sequence:

1. The primary address on the loopback interface **lo0** that is not **127.0.0.1** is used.
2. The primary address on the primary interface is used.

Configuring the Primary Interface for the Router

The *primary interface* for the router has the following characteristics:

- It is the interface that packets go out when you type a command such as ping 255.255.255.255—that is, a command that does not include an interface name (there is no interface ***type*-0/0/0.0** qualifier) and where the destination address does not imply any particular outgoing interface.
- It is the interface on which multicast applications running locally on the router, such as Session Announcement Protocol (SAP), do group joins by default.
- It is the interface from which the default local address is derived for packets sourced out an unnumbered interface if there are no non-127 addresses configured on the loopback interface, lo0.

By default, the multicast-capable interface with the lowest-index address is chosen as the primary interface. If there is no such interface, the point-to-point interface with the lowest index address is chosen. Otherwise, any interface with an address could be picked. In practice, this means that, on the router, the **fxp0** or **em0** interface is picked by default.

To configure a different interface to be the primary interface, include the **primary** statement:

```
primary;
```

You can include this statement at the following hierarchy levels:

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

Configuring the Primary Address for an Interface

The *primary address* on an interface is the address that is used by default as the local address for broadcast and multicast packets sourced locally and sent out the interface. For example, the local address in the packets sent by a **ping interface so-0/0/0.0 255.255.255.255** command is the primary address on interface **so-0/0/0.0**. The primary address flag also can be useful for selecting the local address used for packets sent out unnumbered interfaces when multiple non-127 addresses are configured on the loopback interface, **lo0**. By default, the primary address on an interface is selected as the numerically lowest local address configured on the interface.

To set a different primary address, include the **primary** statement:

```
primary;
```

 Copyright © 2018, Juniper Networks, Inc.

You can include this statement at the following hierarchy levels:

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

Configuring the Preferred Address for an Interface

The *preferred address* on an interface is the default local address used for packets sourced by the local router to destinations on the subnet. By default, the numerically lowest local address is chosen. For example, if the addresses **172.16.1.1/12**, **172.16.1.2/12**, and **172.16.1.3/12** are configured on the same interface, the preferred address on the subnet (by default, **172.16.1.1**) would be used as a local address when you issue a **ping 172.16.1.5** command.

To set a different preferred address for the subnet, include the **preferred** statement:

```
preferred;
```

You can include this statement at the following hierarchy levels:

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

See Also

- Configuring the Interface Address on page 207
- *Junos OS Administration Library*

Operational Behavior of Interfaces When the Same IPv4 Address Is Assigned to Them

You can configure the same IPv4 address on multiple physical interfaces. When you assign the same IPv4 address to multiple physical interfaces, the operational behavior of those interfaces differs, depending on whether they are implicitly or explicitly point-to-point .

NOTE: By default, all interfaces are assumed to be point-to-point (PPP) interfaces. For all interfaces except aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet, you can explicitly configure an interface to be a point-to-point connection.

Copyright © 2018, Juniper Networks, Inc.

NOTE: If you configure the same IP address on multiple interfaces in the same routing instance, Junos OS uses only the first configuration. The remaining IP address configurations are ignored, leaving some interfaces without an assigned address. Interfaces without an assigned address cannot be used as a donor interface for an unnumbered Ethernet interface.

In the following example, the IP address configuration for interface xe-0/0/1.0 is ignored:

```
interfaces {
  xe-0/0/0 {
    unit 0 {
      family inet {
        address 192.168.1.1/24;
      }
    }
  }

}
  xe-0/0/1 {
    unit 0 {
     family inet {
       address 192.168.1.1/24;
     }
  }
}
```

The following examples show the sample configuration of assigning the same IPv4 address to implicitly and explicitly point-to-point interfaces, and their corresponding **show interfaces terse** command outputs to see their operational status.

Configuring same IPv4 address on implicitly PPP interfaces:

```
[edit]
user@host# show
ge-0/1/0 {
  unit 0 {
    family inet {
      address 200.1.1.1/24;
    }
  }
}

ge-3/0/1 {
  unit 0 {
    family inet {
      address 200.1.1.1/24;
    }
  }
}
```

The sample output shown below for the above configuration reveals that only **ge-0/1/0.0** was assigned the same IPv4 address **200.1.1.1/24** and its **link** state was **up**, while **ge-3/0/1.0** was not assigned the IPv4 address, though its **link** state was up,

 Copyright © 2018, Juniper Networks, Inc.

which means that it will be operational only when it gets a unique IPv4 address other than **200.1.1.1/24**.

```
user@host> show interfaces terse ge*
Interface               Admin Link Proto    Local                 Remote
        ge-0/1/0                up    up
        ge-0/1/0.0              up    up   inet      200.1.1.1/24
                                   multiservice
        ge-0/1/1                up    down
        ge-3/0/0                up    down
        ge-3/0/1                up    up
        ge-3/0/1.0              up    up   inet
                                   multiservice
```

Configuring same IPv4 address on explicitly PPP interfaces:

```
[edit]
user@host# show
so-0/0/0 {
  unit 0 {
    family inet {
      address 200.1.1.1/24;
    }
  }
}
so-0/0/3 {
  unit 0 {
    family inet {
      address 200.1.1.1/24;
    }
  }
}
```

The sample output shown below for the above configuration reveals that both **so-0/0/0.0** and **so-0/0/3.0** were assigned the same IPv4 address **200.1.1.1/24** and that their **link** states were down, which means that to make them operational atleast one of them will have be configured with a unique IPv4 address other than **200.1.1.1/24**.

```
user@host> show interfaces terse so*
Interface               Admin Link Proto    Local                 Remote
so-0/0/0                up    up
so-0/0/0.0              up    down inet     200.1.1.1/24
so-0/0/1                up    up
so-0/0/2                up    down
so-0/0/3                up    up
so-0/0/3.0              up    down inet     200.1.1.1/24
so-1/1/0                up    down
so-1/1/1                up    down
so-1/1/2                up    up
so-1/1/3                up    up
so-2/0/0                up    up
so-2/0/1                up    up
so-2/0/2                up    up
so-2/0/3                up    down
```

Related Documentation

- Configuring IPCP Options for Interfaces with PPP Encapsulation on page 216
- Configuring Default, Primary, and Preferred Addresses and Interfaces on page 209

Configuring ICCP for MC-LAG

For multichassis link aggregation (MC-LAG), you must configure Inter-Control Center Communications Protocol (ICCP) to exchange information between two MC-LAG peers.

To enable ICCP, include the **iccp** statement at the **[edit protocols]** hierarchy level:

```
[edit protocols]
iccp {
   authentication-key string;
   local-ip-addr ipv4-address;
   peer ip-address{
      authentication-key string;
      liveness-detection {
         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);
      }
      local-ip-addr ipv4-address;
      redundancy-group-id-list [ redundancy-groups ];
      session-establishment-hold-time value;
   }
   session-establishment-hold-time value;
   traceoptions;
}
```

The **local-ip-address** statement sets the source address. This could be a specified address or interface address. The **session-establishment-hold-time** statement determines whether a chassis takes over as the master at the ICCP session.

The **authentication-key** statement is provided by TCP Message Digest 5 (md5) option for an ICCP TCP session. The **redundancy-group-id-list** statement specifies the redundancy groups between ICCP peers and the **liveness-detection** hierarchy configures Bidirectional Forwarding Detection (BFD) protocol options.

Copyright © 2018, Juniper Networks, Inc.

NOTE: ICCP is based on TCP and it uses IP routes to reach the MC-LAG peer. To ensure that the ICCP session is as resilient as possible, we recommend that you configure alternative routes between the ICCP end-point IP addresses. Alternatively, configure a LAG interface that has two or more interfaces between the MC-LAG pairs to prevent session failure when there are no alternative routes.

For Inter-Control Center Communications Protocol (ICCP) in a multichassis link aggregation group (MC-LAG) configured in an active-active bridge domain, you must ensure that you configure the same peer IP address hosting the MC-LAG by including the **peer *ip-address*** statement at the **[edit protocols iccp]** hierarchy level and the **multi-chassis-protection peer *ip-address*** statement at the **[edit interfaces interface-name]** hierarchy level. Multichassis protection reduces the configuration at the logical interface level for MX Series routers with multichassis aggregated Ethernet (MC-AE) interfaces. If the ICCP is UP and the interchassis data link (ICL) comes UP, the router configured as standby will bring up the MC-AE interfaces shared with the peer active-active node specified by the **peer** statement.

For example, the following statements illustrate how the same peer IP address can be configured for both the ICCP peer and multichassis protection link:

```
set interfaces ae1 unit 0 multi-chassis-protection 10.255.34.112 interface ae0
set protocols iccp peer 10.255.34.112 redundancy-group-id-list 1
```

Although you can commit an MC-LAG configuration with various parameters defined for it, you can configure multichassis protection between two peers without configuring the ICCP peer address. You can also configure multiple ICCP peers and commit such a configuration.

Related Documentation

Copyright © 2018, Juniper Networks, Inc.

Configuring IPCP Options for Interfaces with PPP Encapsulation

For interfaces with PPP encapsulation, you can configure IPCP to negotiate IP address assignments and to pass network-related information such as Windows Name Service (WINS) and Domain Name System (DNS) servers, as defined in RFC 1877, *PPP Internet Protocol Control Protocol Extensions for Name Server Addresses*.

When you enable a PPP interface, you can configure an IP address, enable the interface to negotiate an IP address assignment from the remote end, or allow the interface to be unnumbered. You can also assign a destination profile to the remote end. The destination profile includes PPP properties, such as primary and secondary DNS and NetBIOS Name Servers (NBNSs). These options are described in the following sections:

NOTE: The Junos OS does not request name servers from the remote end; the software does, however, send name servers to the remote end if requested.

Before you begin

You must configure the PPP encapsulation on the interface before configuring the IPCP option. On the logical interface, the following PPP encapsulation types are supported:

- **atm-mlppp-llc**
- **atm-ppp-llc**
- **atm-ppp-vc-mux**
- **multilink-ppp**

For more information about PPP encapsulation, see "Configuring Interface Encapsulation on Logical Interfaces" on page 190 and *Configuring ATM Interface Encapsulation*

- To configure an IP address for the interface, include the **address** statement in the configuration. For more information, see "Configuring the Interface Address" on page 207.

  If you include the **address** statement in the configuration, you cannot include the **negotiate-address** or **unnumbered-address** statement in the configuration.

  When you include the **address** statement in the interface configuration, you can assign PPP properties to the remote end.

  NOTE: The option to negotiate an IP address is not allowed in MLFR and MFR encapsulations.

- To enable the interface to obtain an IP address from the remote end, include the **negotiate-address** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet]** hierarchy level.

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

Copyright © 2018, Juniper Networks, Inc.

user@host# set negotiate-address

NOTE: If you include the **negotiate-address** statement in the configuration, you cannot include the **address** or **unnumbered-address** statement in the configuration.

- To configure an interface to be unnumbered, include the **unnumbered-address** and **destination** statements in the configuration.

  [edit interfaces *interface-name* unit *logical-unit-number* family inet]
  user@host# set unnumbered-address *interface-name*
  user@host# set destination *address*

NOTE:
- The **unnumbered-address** statement enables the local address to be derived from the specified interface. The interface name must include a logical unit number and must have a configured address (see "Configuring the Interface Address" on page 207). Specify the IP address of the remote interface with the **destination** statement.
- If you include the **unnumbered-address** statement in the configuration, you cannot include the **address** or **negotiate-address** statement in the interface configuration.

- To assign PPP properties to the remote end include the **destination-profile** statement:

  [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*]
  user@host# set destination-profile *name*

  [edit interfaces *interface-name* unit *logical-unit-number* family inet unnumbered-address *interface-name*]
  user@host# set destination-profile *name*

NOTE:
- You can assign PPP properties to the remote end, after you include the **address** or **unnumbered-address** statement in the interface configuration.
- You define the profile at the **[edit access group-profile name ppp]** hierarchy level. For more information, see *Example: Group Profile Configuration*

Related Documentation

- *Example: Group Profile Configuration*
- Configuring Interface Encapsulation on Logical Interfaces on page 190

Configuring an Unnumbered Interface

This topic includes the following information:

Overview of Unnumbered Interfaces

When you need to conserve IP addresses, you can configure unnumbered interfaces. Setting up an unnumbered interface enables IP processing on the interface without assigning an explicit IP address to the interface. For IPv6, in which conserving addresses is not a major concern, you can configure unnumbered interfaces to share the same subnet across multiple interfaces. IPv6 unnumbered interfaces are only supported on Ethernet interfaces. The statements you use to configure an unnumbered interface depend on the type of interface you are configuring: a point-to-point interface or an Ethernet interface:

Configuring an Unnumbered Point-to-Point Interface

1. In configuration mode, go to the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

   [edit]
   user@host# edit interfaces *interface-name* unit *logical-unit-number*

2. To configure an unnumbered point-to-point interface, configure the protocol family, but do not include the **address** statement.

   [edit interfaces *interface-name* unit *logical-unit-number*]
   user@host# set family

 Copyright © 2018, Juniper Networks, Inc.

NOTE:

- For interfaces with PPP encapsulation, you can configure an unnumbered interface by including the **unnumbered-interface** statement in the configuration. For more information, see "Configuring IPCP Options for Interfaces with PPP Encapsulation" on page 216.
- When configuring unnumbered interfaces, you must ensure that a source address is configured on some interface in the router. This address is the default address. We recommend that you do this by assigning an address to the loopback interface (**lo0**), as described in "Configuring the Loopback Interface" on page 328. If you configure an address (other than a martian) on the **lo0** interface, that address is always the default address, which is preferable because the loopback interface is independent of any physical interfaces and therefore is always accessible.

Configuring an Unnumbered Ethernet or Demux Interface

1. In configuration mode, go to the **[edit interfaces *interface-name* unit *logical-unit-number* family *family-name*]** hierarchy level.

   [edit]
   user@host# edit interfaces *interface-name* unit *logical-unit-number* family *family-name*

2. To configure an unnumbered Ethernet or demultiplexing interface, include the **unnumbered-address** statement in the configuration.

   [edit interfaces *interface-name* unit *logical-unit-number* family *family-name*]
   user@host# set unnumbered-address *interface-name*

3. (Optional) To specify the unnumbered Ethernet interface as the next-hop interface for a configured static route, include the **qualified-next-hop** statement at the **[edit routing-options static route *destination-prefix*]** hierarchy level. This feature enables you to specify independent preferences and metrics for static routes on a next-hop basis.

   [edit routing-options static route *destination-prefix*]
   user@host# set qualified-next-hop (address | interface-name)

Copyright © 2018, Juniper Networks, Inc.

NOTE:

- The **unnumbered-address** statement currently supports configuration of unnumbered demux interfaces only for the IPv4 address family. You can configure unnumbered Ethernet interfaces for both IPv4 and IPv6 address families.
- The interface that you configure to be unnumbered *borrows* an assigned IP address from another interface, and is referred to as the *borrower interface*. The interface from which the IP address is borrowed is referred to as the *donor interface*. In the **unnumbered-address** statement, ***interface-name*** specifies the donor interface. For an unnumbered Ethernet interface, the donor interface can be an Ethernet, ATM, SONET, or loopback interface that has a logical unit number and configured IP address and is not itself an unnumbered interface. For an unnumbered IP demultiplexing interface, the donor interface can be an Ethernet or loopback interface that has a logical unit number and configured IP address and is not itself an unnumbered interface. In addition, for either Ethernet or demux, the donor interface and the borrower interface must be members of the same routing instance and the same logical system.
- When you configure an unnumbered Ethernet or demux interface, the IP address of the donor interface becomes the source address in packets generated by the unnumbered interface.
- You can configure a host route that points to an unnumbered Ethernet or demux interface. For information about host routes, see the *MPLS Applications Feature Guide*.

Configuring a Preferred Source Address for Unnumbered Ethernet or Demux Interfaces

When a loopback interface with multiple secondary IP addresses is configured as the donor interface for an unnumbered Ethernet or demux interface, you can optionally specify any one of the loopback interface's secondary addresses as the preferred source address for the unnumbered Ethernet or demux interface. This feature enables you to use an IP address other than the primary IP address on some of the unnumbered Ethernet or demux interfaces in your network.

1. In configuration mode, go to the **[edit interfaces *interface-name* unit *logical-unit-number* family *family-name*]** hierarchy level.

   ```
   [edit ]
   user@host# edit interfaces interface-name unit logical-unit-number family family-name
   ```

2. To configure a secondary address on a loopback donor interface as the preferred source address for an unnumbered Ethernet or demux interface, include the **preferred-source-address** option in the **unnumbered-address** statement:

   ```
   [edit interfaces interface-name unit logical-unit-number family family-name]
   user@host# set unnumbered-address interface-name <preferred-source-address
     address
   ```


Copyright © 2018, Juniper Networks, Inc.

NOTE:

The following considerations apply when you configure a preferred source address on an unnumbered Ethernet or demux interface:

- The **unnumbered-address** statement currently supports the configuration of a preferred source address only for the IPv4 address family for demux interfaces, and for IPv4 and IPv6 address families for Ethernet interfaces.
- If you do not specify the preferred source address, the router uses the default primary IP address of the donor interface.
- You cannot delete an address on a donor loopback interface while it is being used as the preferred source address for an unnumbered Ethernet or demux interface.

Restrictions for Configuring Unnumbered Ethernet Interfaces

The following restrictions apply when you configure unnumbered Ethernet interfaces:

- The **unnumbered-address** statement currently supports the configuration of unnumbered Ethernet interfaces for IPv4 and IPv6 address families.
- You cannot assign an IP address to an Ethernet interface that is already configured as an unnumbered interface.
- The donor interface for an unnumbered Ethernet interface must have one or more configured IP addresses.
- The donor interface for an unnumbered Ethernet interfaced cannot be configured as unnumbered.
- An unnumbered Ethernet interface does not support configuration of the following **address** statement options: **arp**, **broadcast**, **primary**, **preferred**, and **vrrp-group**. For information about these options, see "Configuring the Interface Address" on page 207.
- Running IGMP and PIM are supported only on unnumbered Ethernet interfaces that directly face the host and have no downstream PIM neighbors. IGMP and PIM are not supported on unnumbered Ethernet interfaces that act as upstream interfaces in a PIM topology.
- Running OSPF and IS-IS on unnumbered Ethernet interfaces is not supported. However, you can run OSPF over unnumbered Ethernet interfaces configured as a Point-to-Point connection.

Copyright © 2018, Juniper Networks, Inc.

NOTE: If you configure the same address on multiple interfaces in the same routing instance, Junos OS uses only the first configuration, the remaining address configurations are ignored and can leave interfaces without an address. Interfaces that do not have an assigned address cannot be used as a donor interface for an unnumbered Ethernet interface.

For example, in the following configuration the address configuration of interface xe-0/0/1.0 is ignored:

```
interfaces {
  xe-0/0/0 {
    unit 0 {
      family inet {
        address 192.168.1.1/24;
      }
    }
  }
  xe-0/0/1 {
    unit 0 {
     family inet {
       address 192.168.1.1/24;
     }
  }
}
```

For more information on configuring the same address on multiple interfaces, see "Configuring the Interface Address" on page 207.

Displaying the Unnumbered Ethernet Interface Configuration

Purpose To display the configured unnumbered interface at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level:

- Unnumbered interface —ge-1/0/0
- Donor interface —ge-0/0/0
- Donor interface address —4.4.4.1/24

The unnumbered interface "borrows" an IP address from the donor interface.

Action

- Run the **show** command at the **[edit]** hierarchy level.

```
interfaces {
  ge-0/0/0 {
    unit 0 {
      family inet {
        address 4.4.4.1/24;
      }
    }
  }
  ge-1/0/0 {
    unit 0 {
```

 Copyright © 2018, Juniper Networks, Inc.

```
            family inet {
                unnumbered-address ge-0/0/0.0;
            }
        }
    }
}
```

Meaning The sample configuration that is described works correctly on M and T Series routers. For unnumbered interfaces on MX Series routers, you must additionally configure static routes on an unnumbered Ethernet interface by including the **qualified-next-hop** statement at the **[edit routing-options static route destination-prefix]** hierarchy level to specify the unnumbered Ethernet interface as the next-hop interface for a configured static route.

Displaying the Configured Preferred Source Address for an Unnumbered Ethernet Interface

Purpose To display the configuration of preferred source address for an unnumbered interface at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet]** hierarchy level:

- Unnumbered interface —ge-4/0/0
- Donor interface —lo0
- Donor interface primary address—2.2.2.1/32
- Donor interface secondary address—3.3.3.1/32

Action

- Run the **show** command at the **[edit]** hierarchy level.

```
interfaces {
    lo0 {
        unit 0 {
            family inet {
                address 2.2.2.1/32;
                address 3.3.3.1/32;
            }
        }
    }
}
interfaces {
    ge-4/0/0 {
        unit 0 {
            family inet {
                unnumbered-address lo0.0 preferred-source-address 3.3.3.1;
            }
        }
    }
}
```

Meaning The loopback interface **lo0** is the donor interface from which unnumbered Ethernet interface **ge-4/0/0** “borrows” an IP address.

Copyright © 2018, Juniper Networks, Inc.

The example shows one of the loopback interface's secondary addresses, 3.3.3.1, as the preferred source address for the unnumbered Ethernet interface.

Displaying the Configuration for Unnumbered Ethernet Interface as the Next Hop for a Static Route

Purpose To display the unnumbered interface configured as the next hop for the static route at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet]** hierarchy level:

- Unnumbered interface —ge-0/0/0
- Donor interface —lo0
- Donor interface primary address—5.5.5.1/32
- Donor interface secondary address—6.6.6.1/32
- Static route—7.7.7.1/32

Action

- Run the **show** command at the **[edit]** hierarchy level.

```
interfaces {
  lo0 {
    unit 0 {
      family inet {
        address 5.5.5.1/32;
        address 6.6.6.1/32;
      }
    }
  }
}
```

Meaning In this example, **ge-0/0/0** is the unnumbered interface and a loopback interface, **lo0**, is the donor interface from which **ge-0/0/0** "borrows" an IP address. The example also configures a static route to **7.7.7.1/32** with a next hop through unnumbered interface **ge-0/0/0.0**.

Setting the Protocol MTU

When you initially configure an interface, the protocol maximum transmission unit (MTU) is calculated automatically. If you subsequently change the media MTU, the protocol MTU on existing address families automatically changes.

For a list of default protocol MTU values, see "Media MTU Sizes by Interface Type" on page 101.

To modify the MTU for a particular protocol family, include the **mtu** statement:

```
mtu bytes;
```

 Copyright © 2018, Juniper Networks, Inc.

You can include this statement at the following hierarchy levels:

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

If you increase the size of the protocol MTU, you must ensure that the size of the media MTU is equal to or greater than the sum of the protocol MTU and the encapsulation overhead. For a list of encapsulation overhead values, see "Encapsulation Overhead by Interface Encapsulation Type" on page 112. If you reduce the media MTU size, but there are already one or more address families configured and active on the interface, you must also reduce the protocol MTU size. (You configure the media MTU by including the **mtu** statement at the **[edit interfaces *interface-name*]** hierarchy level.)

NOTE: **Changing the media MTU or protocol MTU causes an interface to be deleted and added again.**

The maximum number of data-link connection identifiers (DLCIs) is determined by the MTU on the interface. If you have keepalives enabled, the maximum number of DLCIs is 1000, with the MTU set to 5012.

The actual frames transmitted also contain cyclic redundancy check (CRC) bits, which are not part of the MTU. For example, the default protocol MTU for a Gigabit Ethernet interface is 1500 bytes, but the largest possible frame size is actually 1504 bytes; you need to consider the extra bits in calculations of MTUs for interoperability.

Related Documentation

- Media MTU Overview on page 100
- Configuring the Media MTU on page 115

Disabling the Removal of Address and Control Bytes

For Point-to-Point Protocol (PPP) CCC-encapsulated interfaces, the address and control bytes are removed by default before the packet is encapsulated into a tunnel.

You can disable the removal of address and control bytes. To do this, include the **keep-address-and-control** statement:

```
keep-address-and-control;
```

You can include this statement at the following hierarchy levels:

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

Related Documentation

- keep-address-and-control on page 709

Copyright © 2018, Juniper Networks, Inc.

Disabling the Transmission of Redirect Messages on an Interface

By default, the interface sends protocol redirect messages. To disable the sending of these messages on an interface, include the **no-redirects** statement:

```
no-redirects;
```

You can include this statement at the following hierarchy levels:

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

To disable the sending of protocol redirect messages for the entire router or switch, include the **no-redirects** statement at the **[edit system]** hierarchy level.

Related Documentation

- no-redirects on page 825

Applying Policers

Overview of Applying Policers

Policers allow you to perform simple traffic policing on specific interfaces or Layer 2 virtual private networks (VPNs) without configuring a firewall filter. To apply policers, include the **policer** statement:

```
policer {
    arp policer-template-name;
    input policer-template-name;
    output policer-template-name;
}
```

You can include these statements at the following hierarchy levels:

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

In the **family** statement, the protocol family can be **ccc**, **inet**, **inet6**, **mpls**, **tcc**, or **vpls**.

Copyright © 2018, Juniper Networks, Inc.

In the **arp** statement, list the name of one policer template to be evaluated when Address Resolution Protocol (ARP) packets are received on the interface. By default, an ARP policer is installed that is shared among all the Ethernet interfaces on which you have configured the **family inet** statement. If you want more stringent or lenient policing of ARP packets, you can configure an interface-specific policer and apply it to the interface. You configure an ARP policer just as you would configure any other policer, at the **[edit firewall policer]** hierarchy level. If you apply this policer to an interface, the default ARP packet policer is overridden. If you delete this policer, the default policer takes effect again.

In the **input** statement, list the name of one policer template to be evaluated when packets are received on the interface.

In the **output** statement, list the name of one policer template to be evaluated when packets are transmitted on the interface.

NOTE: **To use policing on a CCC or TCC interface, you must configure the CCC or TCC protocol family.**

You can configure a different policer on each protocol family on an interface, with one input policer and one output policer for each family. When you apply policers, you can configure the family **ccc**, **inet**, **inet6**, **mpls**, **tcc**, or **vpls** only, and one ARP policer for the family **inet** protocol only. Each time a policer is referenced, a separate copy of the policer is installed on the packet forwarding components for that interface.

If you apply both policers and firewall filters to an interface, input policers are evaluated before input firewall filters, and output policers are evaluated after output firewall filters.

If you apply the policer to the interface **lo0**, it is applied to packets received or transmitted by the Routing Engine.

On T Series, M120, and M320 platforms, if the interfaces are on the same FPC, the filters or policers do not act on the sum of traffic entering and exiting the interfaces.

For more information about policers, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*.

Applying Aggregate Policers

Applying Aggregate Policers

By default, if you apply a policer to multiple protocol families on the same logical interface, the policer restricts traffic for each protocol family individually. For example, a policer with a 50 Mbps bandwidth limit applied to both IPv4 and IPv6 traffic would allow the interface to accept 50 Mbps of IPv4 traffic and 50 Mbps of IPv6 traffic. If you apply an aggregate policer, the policer would allow the interface to receive only 50 Mbps of IPv4 and IPv6 traffic combined.

Copyright © 2018, Juniper Networks, Inc.

To configure an aggregate policer, include the **logical-interface-policer** statement at the **[edit firewall policer *policer-template-name*]** hierarchy level:

```
[edit firewall policer policer-template-name]
logical-interface-policer;
```

For the policer to be treated as an aggregate, you must apply it to multiple protocol families on a single logical interface by including the **policer** statement:

```
policer {
   arp policer-template-name;
   input policer-template-name;
   output policer-template-name;
}
```

You can include these statements at the following hierarchy levels:

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

In the **family** statement, the protocol family can be **ccc**, **inet**, **inet6**, **mpls**, **tcc**, or **vpls**.

The protocol families on which you do not apply the policer are not affected by the policer. For example, if you configure a single logical interface to accept MPLS, IPv4, and IPv6 traffic and you apply the logical interface policer **policer1** to only the IPv4 and IPv6 protocol families, MPLS traffic is not subject to the constraints of **policer1**.

If you apply **policer1** to a different logical interface, there are two instances of the policer. This means the Junos OS polices traffic on separate logical interfaces separately, not as an aggregate, even if the same logical-interface policer is applied to multiple logical interfaces on the same physical interface port.

Example: Applying Aggregate Policers

Configure two logical interface policers: **aggregate_police1** and **aggregate_police2**. Apply **aggregate_police1** to IPv4 and IPv6 traffic received on logical interface **fe-0/0/0.0**. Apply **aggregate_police2** to CCC and MPLS traffic received on logical interface fe-0/0/0.0. This configuration causes the software to create only one instance of **aggregate_police1** and one instance of **aggregate_police2**.

Apply **aggregate_police1** to IPv4 and IPv6 traffic received on another logical interface **fe-0/0/0.1**. This configuration causes the software to create a new instance of **aggregate_police1**, one that applies to unit 0 and another that applies to unit 1.

```
[edit firewall]
policer aggregate_police1 {
   logical-interface-policer;
   if-exceeding {
      bandwidth-limit 100m;
      burst-size-limit 500k;
   }
   then {
      discard;
```

 Copyright © 2018, Juniper Networks, Inc.

```
    }
}
policer aggregate_police2 {
    logical-interface-policer;
    if-exceeding {
        bandwidth-limit 10m;
        burst-size-limit 200k;
    }
    then {
        discard;
    }
}
[edit interfaces fe-0/0/0]
unit 0 {
    family inet {
        policer {
            input aggregate_police1;
        }
    }
    family inet6 {
        policer {
            input aggregate_police1;
        }
    }
    family ccc {
        policer {
            input aggregate_police2;
        }
    }
    family mpls {
        policer {
            input aggregate_police2;
        }
    }
}
unit 1 {
    family inet {
        policer {
            input aggregate_police1;
        }
    }
    family inet6 {
        policer {
            input aggregate_police1;
        }
    }
}
```

Applying Hierarchical Policers on Enhanced Intelligent Queuing PICs

• Applying Hierarchical Policers on Enhanced Intelligent Queuing PICs on page 229

Applying Hierarchical Policers on Enhanced Intelligent Queuing PICs

M40e, M120, and M320 edge routers and T Series core routers with Enhanced Intelligent Queuing (IQE) PICs support hierarchical policers in the ingress direction and allow you to apply a hierarchical policer for the premium and aggregate (premium plus normal) traffic levels to an interface. Hierarchical policers provide cross-functionality between the configured physical interface and the Packet Forwarding Engine.

Before you begin, there are some general restrictions that apply to hierarchical policers:

- Only one type of policer can be configured for a logical or physical interface. For example, a hierarchical policer and a regular policer in the same direction for the same logical interface is not allowed.
- The chaining of the policers—that is, applying policers to both a port and the logical interfaces of that port—is not allowed.
- There is a limit of 64 policers per interface in case there is no BA classification, providing a single policer per DLCI.
- Only one kind of policer can be applied on a physical or logical interface.
- The policer should be independent of BA classification. Without BA classification, all traffic on an interface will be treated either as EF or non-EF, based on the configuration. With BA classification, an interface can support up to 64 policers. Again, the interface here may be a physical interface or logical interface (for example, DLCI).
- With BA classification, the miscellaneous traffic (the traffic *not* matching with any of the BA classification DSCP/EXP bits) will be policed as non-EF traffic. No separate policers will be installed for this traffic.

Hierarchical Policer Overview

Hierarchical policing uses two token buckets, one for aggregate (non-EF) traffic and one for premium (EF) traffic. Which traffic is EF and which is non-EF is determined by the class-of-service configuration. Logically, hierarchical policing is achieved by chaining two policers.

Figure 11: Hierarchical Policer

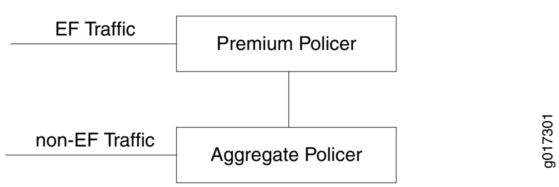

In the example in Figure 11 on page 230, EF traffic is policed by Premium Policer and non EF traffic is policed by Aggregate Policer. What that means is, for EF traffic the out-of-spec action will be the one that is configured for Premium Policer, but the in-spec EF traffic will still consume the tokens from the Aggregate Policer.

But EF traffic will never be submitted to the out-of-spec action of the Aggregate Policer. Also, if the out-of-spec action of the Premium Policer is not set to Discard, those out-of-spec packets will not consume the tokens from the Aggregate Policer. Aggregate Policer only polices the non-EF traffic. As you can see, the Aggregate Policer token bucket can go negative, if all the tokens are consumed by the non-EF traffic and then you get bursts of EF traffic. But that will be for a very short time, and over a period of time it will average out. For example:

 Copyright © 2018, Juniper Networks, Inc.

- *Premium Policer:* Bandwidth 2 Mbps, OOS Action: Discard
- *Aggregate Policer:* Bandwidth 10 Mbps, OOS Action: Discard

In the above case, EF traffic is guaranteed 2 Mbps and the non-EF traffic will get from 8 Mbps to 10 Mbps, depending on the input rate of the EF traffic.

Hierarchical Policing Characteristics

Hierarchical token bucket features include:

- Ingress traffic is first classified into EF and non-EF traffic prior to applying a policer:
  - Classification is performed by Q-tree lookup
- Channel number selects a shared token bucket policer:
  - Dual token bucket policer is divided into two single bucket policers:
    - Policer1—EF traffic
    - Policer2—non-EF traffic
- Shared token bucket is used to police the traffic as follows:
  - Policer1 is set to EF rate (for example, 2 Mbps)
  - Policer2 is set to aggregate interface policed rate (for example, 10 Mbps).
  - EF traffic gets applied to Policer1.
    - If traffic is in-spec it is allowed to pass and decrement from both Policer1 and Policer2.
    - If traffic is out-of-spec it can be discarded or marked with a new FC or loss priority. Policer2 will not do anything with out-of-spec EF traffic.
  - Non-EF traffic gets applied only to Policer2.
    - If traffic is in-spec it is allowed to pass through and decremented Policer2.
    - If traffic is out-of-spec it is discarded or marked with a new FC or set with a new drop priority.
- Rate-limit the port speed to a desired rate at Layer 2
- Rate-limit the EF traffic
- Rate-limit the non-EF traffic
- Policing drops counted per color

See Also

- *Class of Service Feature Guide for Routing Devices and EX9200 Switches*

Configuring Hierarchical Policers

To configure a hierarchical policer, apply the **policing-priority** statement to the proper forwarding class and configure a hierarchical policer for the aggregate and premium level. For more information about class of service, see the *Class of Service Feature Guide for Routing Devices and EX9200 Switches*.

NOTE: Hierarchical policers can only be configured on SONET physical interfaces hosted on an IQE PIC. Only aggregate and premium levels are supported.

CoS Configuration of Forwarding Classes for Hierarchical Policers

```
[edit class-of-service forwarding-classes]
class fc1 queue-num 0 priority high policing-priority premium;
class fc2 queue-num 1 priority low policing-priority normal;
class fc3 queue-num 2 priority low policing-priority normal;
class fc4 queue-num 3 priority low policing-priority normal;
```

For detailed information on class-of-service configuration and statements, see the *Class of Service Feature Guide for Routing Devices and EX9200 Switches*.

Firewall Configuration for Hierarchical Policers

```
[edit firewall hierarchical-policer foo]
aggregate {
  if-exceeding {
    bandwidth-limit 70m;
    burst-size-limit 1500;
  }
  then {
  discard;
}
premium {
  if-exceeding {
    bandwidth-limit 50m;
    burst-size-limit 1500;
  }
  then {
    discard;
  }
}
```

You can apply the hierarchical policer as follows:

```
[edit interfaces so-0/1/0 unit 0 layer2-policer]
input-hierarchical-policer foo;
```

You also have the option to apply the policer at the physical port level as follows:

```
[edit interfaces so-0/1/0 layer2-policer]
input-hierarchical-policer foo;
```

Configuring a Single-Rate Two-Color Policer

You can configure a single-rate two-color policer as follows:

```
[edit firewall policer foo]
  if-exceeding {
    bandwidth-limit 50m;
    burst-size-limit 1500;
  }
```


Copyright © 2018, Juniper Networks, Inc.

```
then {
   discard;
}
```

You can apply the policer as follows:

```
[edit interfaces so-0/1/0 unit 0 layer2-policer]
input-policer foo;
```

You also have the option to apply the policer at the physical port level as follows:

```
[edit interfaces so-0/1/0 layer2-policer]
input-policer foo;
```

Configuring a Single-Rate Color-Blind Policer

This section describes single-rate color blind and color aware policers.

You can configure a single-rate color blind policer as follows:

```
[edit firewall three-color-policer foo]
single-rate {
   color-blind;
   committed-information-rate 50m;
   committed-burst-size 1500;
   excess-burst-size 1500;
}
```

You can apply the single-rate color blind policer as follows:

```
[edit interfaces so-0/1/0 unit 0 layer2-policer]
input-three-color foo;
```

You can configure a single-rate color-aware policer as follows:

```
[edit firewall three-color-policer bar]
single-rate {
   color-aware;
   committed-information-rate 50m;
   committed-burst-size 1500;
   excess-burst-size 1500;
}
```

You can apply the single-rate color-aware policer as follows:

```
[edit interfaces so-0/1/0 unit 0 layer2-policer]
input-three-color foo;
```

You also have the option to apply the policer at the physical port level as follows:

```
[edit interfaces so-0/1/0 layer2-policer]
input-three-color bar;
```

Configuring a Two-Rate Tricolor Marker Policer

Ingress policing is implemented using a two-rate tricolor marker (trTCM). This is done with a dual token bucket (DTB) that maintains two rates, committed, and a peak. Egress static policing also uses a token bucket.

Copyright © 2018, Juniper Networks, Inc.

The token buckets perform the following ingress policing functions:

- (1K) trTCM - Dual token bucket (red, yellow, and green marking)
- Policing is based on Layer 2 packet size:
  - After +/- byte adjust offset
- Marking is color aware and color blind:
  - Color aware needs to have the color set by q-tree lookup based on:
    - ToS
    - EXP
- Programmable marking actions:
  - Color (red, yellow, green)
  - Drop based on color and congestion profile
- Policer is selected based on the arriving channel number:
  - Channel number LUT produces policer index and queue index
  - Multiple channels can share the same policer (LUT produces same policer index)
- Support ingress policing and trTCM at the following levels:
  - Queue
  - Logical interface (ifl/DLCI)
  - Physical interface (ifd)
  - Physical port (controller ifd)
  - Any combinations of logical interface, physical interface, and port
- Support percentage of interface speed and bits per second

Rate limits may be applied to selected queues on ingress and on predefined queues at egress. The token bucket operates in color aware and color blind modes (specified by RFC 2698).

Configuring a Color-Blind trTCM

```
[edit firewall three-color-policer foo]
two-rate {
  color-blind;
  committed-information-rate 50m;
  committed-burst-size 1500;
  peak-information-rate 100m;
  peak-burst-size 3k;
}
```

You can apply the three-color two-rate color-blind policer as follows:

```
[edit interfaces so-0/1/0 unit 0 layer2-policer]
input-three-color foo;
```

 Copyright © 2018, Juniper Networks, Inc.

You also have the option to apply the policer at the physical port level as follows:

```
[edit interfaces so-0/1/0 layer2-policer]
input-three-color foo;
```

Configuring a Color-Aware trTCM

```
[edit firewall three-color-policer bar]
two-rate {
  color-aware;
  committed-information-rate 50m;
  committed-burst-size 1500;
  peak-information-rate 100m;
  peak-burst-size 3k;
}
```

You can apply the three-color two-rate color-aware policer as follows:

```
[edit interfaces so-0/1/0 unit 0 layer2-policer]
input-three-color bar;
```

You also have the option to apply the policer at the physical port level as follows:

```
[edit interfaces so-0/1/0 layer2-policer]
input-three-color bar;
```

See Also

- *Class of Service Feature Guide for Routing Devices and EX9200 Switches*

Applying a Filter to an Interface

Defining Interface Groups in Firewall Filters

When applying a firewall filter, you can define an interface to be part of an *interface group*. Packets received on that interface are tagged as being part of the group. You can then match these packets using the **interface-group** match statement, as described in the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*.

To define the interface to be part of an interface group, include the **group** statement:

group *filter-group-number*;

You can include this statement at the following hierarchy levels:

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

NOTE: The number 0 is not a valid interface group number.

Filter-Based Forwarding on the Output Interface

If port-mirrored packets are to be distributed to multiple monitoring or collection interfaces, based on patterns in packet headers, it is helpful to configure a filter-based forwarding (FBF) filter on the port-mirroring egress interface.

When an FBF filter is installed as an output filter, a packet that is forwarded to the filter has already undergone at least one route lookup. After the packet is classified at the egress interface by the FBF filter, it is redirected to another routing table for additional route lookup. To avoid packet looping inside the Packet Forwarding Engine, the route lookup in the latter routing table (designated by an FBF routing instance) must result in a different next hop from any next hop specified in a table that has already been applied to the packet.

If an input interface is configured for FBF, the source lookup is disabled for those packets headings to a different routing instance, since the routing table is not set up to handle the source lookup.

For more information about FBF configuration, see the *Junos OS Routing Protocols Library*. For more information about port mirroring, see the *Junos OS Services Interfaces Library for Routing Devices*.

Applying a Filter to an Interface

To apply firewall filters to an interface, include the **filter** statement:

```
filter {
  group filter-group-number;
  input filter-name;
  input-list [ filter-names ];
  output filter-name;
  output-list [ filter-names ];
}
```

To apply a single filter, include the **input** statement:

```
filter {
  input filter-name;
}
```

To apply a list of filters to evaluate packets received on an interface, include the **input-list** statement.

```
filter {
  input-list [ filter-names ];
}
```

Up to 16 filter names can be included in an input list.

To apply a list of filters to evaluate packets transmitted on an interface, include the **output-list** statement.

```
filter {
  output-list [ filter-names ];
}
```

 Copyright © 2018, Juniper Networks, Inc.

When you apply filters using the **input-list** statement or the **output-list** statement, a new filter is created with the name <*interface-name*>.<*unit-direction*>. This filter is exclusively interface-specific.

You can include these statements at the following hierarchy levels:

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

In the **family** statement, the protocol family can be **ccc**, **inet**, **inet6**, **mpls**, or **vpls**.

In the **group** statement, specify the interface group number to associate with the filter.

In the **input** statement, list the name of one firewall filter to be evaluated when packets are received on the interface.

In the **input-list** statement, list the names of filters to evaluate when packets are received on the interface. You can include up to 16 filter names.

In the **output** statement, list the name of one firewall filter to be evaluated when packets are transmitted on the interface.

NOTE: Output filters do not work for broadcast and multicast traffic, including VPLS traffic (except in MX Series routers with MPC/MIC interfaces), as shown in "Applying a Filter to an Interface" on page 236.

NOTE: On an MX Series router, you cannot apply as an output filter, a firewall filter configured at the **[edit firewall filter family ccc]** hierarchy level. Firewall filters configured for the **family ccc** statement can be applied only as input filters.

In the **output-list** statement, list the names of filters to evaluate when packets are transmitted on the interface. You can include up to 16 filter names.

You can use the same filter one or more times. On M Series routers (except the M320 and M120 routers), if you apply a firewall filter or policer to multiple interfaces, the filter or policer acts on the sum of traffic entering or exiting those interfaces.

On T Series, M120, and M320 routers, interfaces are distributed among multiple packet forwarding components. Therefore, on these routers, if you apply a firewall filter or policer to multiple interfaces, the filter or policer acts on the traffic stream entering or exiting each interface, regardless of the sum of traffic on the multiple interfaces.

For more information on Understanding Ethernet Frame Statistics, see the *MX Series Layer 2 Configuration Guide*.

Copyright © 2018, Juniper Networks, Inc.

If you apply the filter to the interface **lo0**, it is applied to packets received or transmitted by the Routing Engine. You cannot apply MPLS filters to the management interface (**fxp0** or **em0**) or the loopback interface (**lo0**).

Filters applied at the **[set interfaces lo0 unit 0 family any filter input]** hierarchy level are not installed on T4000 Type 5 FPCs.

For more information about firewall filters, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*. For more information about MPLS filters, see the *MPLS Applications Feature Guide*.

Example: Input Filter for VPLS Traffic

For M Series and T Series routers only, apply an input filter to VPLS traffic. Output filters do not work for broadcast and multicast traffic, including VPLS traffic. Note that on MX Series routers with MPC/MIC interfaces, the VPLS filters on the egress is applicable to broadcast, multicast, and unknown unicast traffic.

```
[edit interfaces]
fe-2/2/3 {
  vlan-tagging;
  encapsulation vlan-vpls;
  unit 601 {
    encapsulation vlan-vpls;
    vlan-id 601;
    family vpls {
      filter {
        input filter1;  # Works for multicast destination MAC address
        output filter1;  # Does not work for multicast destination MAC address
      }
    }
  }
}
[edit firewall]
family vpls {
  filter filter1 {
    term 1 {
      from {
        destination-mac-address {
          01:00:0c:cc:cc:cd/48;
        }
      }
      then {
        discard;
      }
    }
    term 2 {
      then {
        accept;
      }
    }
  }
}
```

 Copyright © 2018, Juniper Networks, Inc.

Example: Filter-Based Forwarding at the Output Interface

The following example illustrates the configuration of filter-based forwarding at the output interface. In this example, the packet flow follows this path:

1. A packet arrives at interface **fe-1/2/0.0** with source and destination addresses **10.50.200.1** and **10.50.100.1** respectively.
2. The route lookup in routing table **inet.0** points to the egress interface **so-0/0/3.0**.
3. The output filter installed at **so-0/0/3.0** redirects the packet to routing table **fbf.inet.0**.
4. The packet matches the entry **10.50.100.0/25** in the **fbf.inet.0** table, and finally leaves the router from interface **so-2/0/0.0**.

```
[edit interfaces]
so-0/0/3 {
  unit 0 {
    family inet {
      filter {
        output fbf;
      }
      address 10.50.10.2/25;
    }
  }
}
fe-1/2/0 {
  unit 0 {
    family inet {
      address 10.50.50.2/25;
    }
  }
}
so-2/0/0 {
  unit 0 {
    family inet {
      address 10.50.20.2/25;
    }
  }
}
[edit firewall]
filter fbf {
  term 0 {
    from {
      source-address {
        10.50.200.0/25;
      }
    }
    then routing-instance fbf;
  }
  term d {
    then count d;
  }
}
[edit routing-instances]
fbf {
  instance-type forwarding;
  routing-options {
```

```
            static {
               route 10.50.100.0/25 next-hop so-2/0/0.0;
            }
         }
      }
      [edit routing-options]
      interface-routes {
         rib-group inet fbf-group;
      }
      static {
         route 10.50.100.0/25 next-hop 10.50.10.1;
      }
      rib-groups {
         fbf-group {
            import-rib [inet.0 fbf.inet.0];
         }
      }
```

Guidelines for Configuring Unicast RPF on ACX Series Routers

Observe the following guidelines while configuring unicast RPF on ACX Series routers:

- Support for physical interfaces impacts inet families only.
- The RPF check to be used when routing is asymmetrical is not supported because the **unicast-reverse-path (active-paths | feasible-paths)** statement at the **[edit routing-instances** ***routing-instance-name*** **instance-type** ***name*** **routing-options forwarding-table]** hierarchy level is not supported.
- Even if uRPF checking is enabled, the reverse path checking is not performed if the following conditions apply:
  - The destination IP address is not a unicast address. This applies for both IPV4 and IPV6 packets.
  - The source IP address is IPV6 and the address is a link local address (FE80::/10)
  - The received packet is a BOOTP/DHCP packet (SIP=0.0.0.0 and DIP=255.255.255.255)
- If you enable/disable unicast RPF on live traffic, some packets are dropped while the packet forwarding components are updating. This behavior occurs because route reinstallation is initiated while you enable or disable uRPF.
- uRPF is supported at the logical interface level. Due to hardware limitations, support is available only at the logical interface level.
- Strict mode on ECMP routes is not supported in ACX. This condition occurs because the hardware treats ECMP routes as Loose Mode although the port is configured as Strict mode. Because ECMP uses multiple physical paths for the route the reverse path check results in utilizing many paths (routes) and the source port validation method is not used in case of Strict mode. As a result, such a network scenario operates in the same manner as loose mode.

Copyright © 2018, Juniper Networks, Inc.

- When the strict mode is enabled on the interface, if the packet is coming with an SIP address which ARP resolution is pending will be dropped as it points to RESOLVE_NH.
- uRPF fail filter can be configured for family <*inet* | *inet6*> in ACX.

NOTE: The uRPF fail filter cannot match packets failed at ingress port check (strict mode).

The uRPF fail filter can match packets failing source IP lookup but cannot match packets failing the input interface check (strict mode).

The uRPF fail filter applies only to interface-specific instances of the firewall filter.

The uRPF fail filters do not support **reject** and **routing-instance** actions.

- uRPF can be configured for family <*inet* | *inet6*> on IRB interfaces in ACX.
- uRPF implementation in ACX does not consider all feasible paths for reverse path verification and only active path based verification is supported.
- uRPF failure packets statistics are not supported in ACX.
- You can use either the **show interfaces extensive** command or the **show interfaces detail** command to verify that unicast RPF is enabled and working on the interface. In the **Flags** section of the output, if unicast reverse-path forwarding (RPF) is explicitly configured on the specified interface, the uRPF flag is displayed. If unicast RPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag is not displayed even though unicast RPF is enabled.
- The uRPF detail in the **Flags** section of the output of the **show interfaces (detail | extensive)** commands is displayed only for logical interfaces on which uRPF is configured. Otherwise, this information is not shown.

Related Documentation

-

Configuring Unicast RPF

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:

- You can protect a network by applying unicast RPF check feature at the edge (on customer facing interfaces) of the network. In an ISP environment, this can impact the network which can impose on a scaled setup. In case if you have already protected the edge of your network, a packet with a spoofed IP source address would not even appear in a core facing interface. In this case, unicast RPF check is not necessary. Enabling unicast RPF feature can impact the control plane performance, so use it where it is required. So it is strongly recommended not to enable this feature on the network core (internal) interfaces.

The following sections describe unicast RPF in detail:

Unicast RPF and Default Route

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 Library*.

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 243
- Unicast RPF Behavior Without a Default Route on page 243
- Unicast RPF with Routing Asymmetry on page 244

 Copyright © 2018, Juniper Networks, Inc.

Unicast RPF Behavior with a Default Route

On all routers except those with MPCs and the MX80 router, unicast RPF behaves as follows if you configure a default route that uses an interface configured with unicast RPF:

- 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 the source address of the packet matches any of the routes (either default or learned) that can be reachable through 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.

On all routers with MPCs and the MX80 router, unicast RPF behaves as follows if you configure a default route that uses an interface configured with unicast RPF:

- Loose mode—All packets except the packets whose source is learned from the default route are accepted. All packets whose source is learned from the default route are dropped at the Packet Forwarding Engine. The default route is treated as if the route does not exist.
- Strict mode—The packet is accepted when the source address of the packet matches any of the routes (either default or learned) that can be reachable through 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.

On all routers, 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 244 and "Configuring Unicast RPF Loose Mode" on page 246. 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.

Copyright © 2018, Juniper Networks, Inc.

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 12 on page 244 shows unicast RPF in an environment with routing asymmetry.

Figure 12: Unicast RPF with Routing Asymmetry

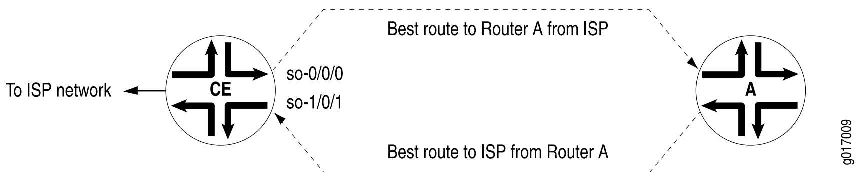

In Figure 12 on page 244, if you enable unicast RPF on interface **so-0/0/0**, traffic destined for Router A is not rejected. If you enable unicast RPF on interface **so-1/0/1**, traffic from Router A is rejected.

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 "Configuring Unicast RPF" on page 250.

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 "Configuring Unicast RPF" on page 250.

For more information about unicast RPF, see the *Junos OS Routing Protocols Library*. For more information about defining fail filters, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*.

 Copyright © 2018, Juniper Networks, Inc.

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.

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**.

To configure unicast RPF in strict mode:

1. Configure the fail filter:

```
[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;
    }
```

```
    }
  }
```

2. Configure unicast RPF on interfaces:

```
[edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family inet {
        rpf-check fail-filter rpf-special-case-dhcp-bootp;
      }
    }
  }
}
```

3. Commit the configuration.

```
[edit]
commit;
```

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**:

1. 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*>]**

2. For example:

   In this example, no special configuration beyond device initialization is required.

   Configure unicast RPF loose 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**.

   To configure unicast RPF in loose mode:

   a. Configure the fail filter:

 Copyright © 2018, Juniper Networks, Inc.

```
[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;
      }
   }
}
```

b. Configure unicast RPF on interfaces:

```
[edit]
interfaces {
   so-0/0/0 {
      unit 0 {
         family inet {
            rpf-check fail-filter rpf-special-case-dhcp-bootp;
            mode loose;
         }
      }
   }
}
```

c. Commit the configuration.

```
[edit]
commit;
```

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;
  }
}
```

In this example, no special configuration beyond device initialization is required.

Configure unicast RPF loose 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**.

To configure unicast RPF loose mode with the ability to discard packets:

1. Configure the fail filter:

```
[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;
    }
  }
}
```

2. Configure unicast RPF on interfaces:

```
[edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family inet {
```

 Copyright © 2018, Juniper Networks, Inc.

```
                rpf-check fail-filter rpf-special-case-dhcp-bootp;
                mode loose;
            }
        }
    }
}
```

3. Configure the ability to discard packets.

```
[edit]
forwarding-options{
  rpf-loose-mode-discard {
    family {
      inet;
      }
    }
  }
```

4. Commit the configuration.

```
[edit]
commit;
```

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 Library for Routing Devices*. For more information about FBF, see the *Junos OS Routing Protocols Library*.

Configure unicast RPF on a Layer 3 VPN interface:

```
[edit interfaces]
so-0/0/0 {
  unit 0 {
    family inet {
      rpf-check;
    }
  }
```

Copyright © 2018, Juniper Networks, Inc.

```
}
[edit routing-instance]
VPN-A {
  interface so-0/0/0.0;
}
```

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;
      }
    }
  }
}
```

See Also

- *unicast-reverse-path*
- Example: Configuring Unicast Reverse-Path-Forwarding Check on page 253

Verifying Unicast RPF Status

Purpose Verify that unicast reverse-path forwarding (RPF) is enabled and is working on the interface.

 Copyright © 2018, Juniper Networks, Inc.

Action Use one of the **show interfaces** ***interface-name*** commands with either the **extensive** or **detail** options to verify that unicast RPF is enabled and working on the switch. The example below displays output from the **show interfaces ge- extensive** command.

```
user@switch> show interfaces ge1/0/10 extensive
Physical interface: ge-1/0/10, Enabled, Physical link is Down
  Interface index: 139, SNMP ifIndex: 58, Generation: 140
  Link-level type: Ethernet, MTU: 1514, Speed: Auto, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
  Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:19:e2:50:95:ab, Hardware address: 00:19:e2:50:95:ab
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                    0 bps
   Output bytes  :                   0                    0 bps
   Input  packets:                   0                    0 pps
   Output packets:                   0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                  0
    Output bytes  :                  0
    Input  packets:                  0
    Output packets:                  0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, 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 assured-forw                  0                    0                    0

    5 expedited-fo                  0                    0                    0

    7 network-cont                  0                    0                    0

  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
    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
    Autonegotiation information:
      Negotiation status: Incomplete
    Packet Forwarding Engine configuration:
      Destination slot: 1

    Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 59) (Generation 135)
      Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
      Traffic statistics:
       Input  bytes  :                    0
       Output bytes  :                    0
       Input  packets:                    0
       Output packets:                    0
       IPv6 transit statistics:
        Input  bytes  :                   0
        Output bytes  :                   0
        Input  packets:                   0
        Output packets:                   0
      Local statistics:
       Input  bytes  :                    0
       Output bytes  :                    0
       Input  packets:                    0
       Output packets:                    0
      Transit statistics:
       Input  bytes  :                    0                    0 bps
       Output bytes  :                    0                    0 bps
       Input  packets:                    0                    0 pps
       Output packets:                    0                    0 pps
       IPv6 transit statistics:
        Input  bytes  :                   0
        Output bytes  :                   0
        Input  packets:                   0
        Output packets:                   0
          Protocol inet, Generation: 144, Route table: 0
        Flags: uRPF
        Addresses, Flags: Is-Preferred Is-Primary
```

Meaning The **show interfaces ge-1/0/10 extensive** command (and the **show interfaces ge-1/0/10 detail** command) displays in-depth information about the interface. The **Flags:** output field near the bottom of the display reports the unicast RPF status. If unicast RPF has not been enabled, the **uRPF** flag is not displayed.

On EX3200 and EX4200 switches, unicast RPF is implicitly enabled on *all* switch interfaces, including aggregated Ethernet interfaces (also referred to as link aggregation groups or LAGs) and routed VLAN interfaces (RVIs) when you enable unicast RPF on a

 Copyright © 2018, Juniper Networks, Inc.

single interface. However, the unicast RPF status is shown as enabled only on interfaces for which you have explicitly configured unicast RPF. Thus, the **uRPF** flag is not displayed on interfaces for which you have not explicitly configured unicast RPF even though unicast RPF is implicitly enabled on all interfaces on EX3200 and EX4200 switches.

Related Documentation

- *show interfaces xe*
- *Example: Configuring Unicast RPF on an EX Series Switch*
- *Configuring Unicast RPF on ACX Series Routers*
- *Configuring Unicast RPF (CLI Procedure)*
- *Disabling Unicast RPF (CLI Procedure)*
- *Troubleshooting Unicast RPF*

Example: Configuring Unicast Reverse-Path-Forwarding Check

Understanding How Unicast Reverse Path Forwarding Prevents Spoofed IP Packet 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.

See Also

- Example: Configuring Unicast Reverse-Path-Forwarding Checking to Prevent DoS and DDoS Attacks on page 254

Example: Configuring Unicast Reverse-Path-Forwarding Checking to Prevent DoS and DDoS Attacks

Unicast reverse path forwarding (RPF) helps protect against DoS and DDoS attacks by verifying the unicast source address of each packet that arrives on an ingress interface where unicast RPF is enabled.

This example shows how to help defend ingress interfaces against denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks by configuring unicast RPF to filter incoming traffic.

Requirements

In this example, no special configuration beyond device initialization is required.

Overview

Large amounts of unauthorized traffic such as attempts to flood a network with fake (bogus) service requests in a DoS attack can consume network resources and deny service to legitimate users. One way to help prevent DoS and DDoS attacks is to verify that incoming traffic originates from legitimate network sources.

Unicast RPF helps ensure that a traffic source is legitimate (authorized) by comparing the source address of each packet that arrives on an interface to the forwarding table entry for its source address. If the device uses the same interface that the packet arrived on to reply to the packet's source, this verifies that the packet originated from an authorized source, and the device forwards the packet. If the device does not use the same interface that the packet arrived on to reply to the packet's source, the packet might have originated from an unauthorized source, and the device discards the packet.

In this example, Device B has unicast RPF configured. Device A is using OSPF to advertise a prefix for the link that connects to Device D. OSPF is enabled on the links between Device B and Device C and the links between Device A and Device C, but not on the links between Device A and Device B. Therefore, Device B learns about the route to Device D through Device C.

If ingress filtering is used in an environment where DHCP or BOOTP is used, it should be ensured that the packets with a source address of 0.0.0.0 and a destination address of 255.255.255.255 are allowed to reach the relay agent in routers when appropriate.

This example also includes a fail filter. When a packet fails the unicast RPF check, the fail filter is evaluated to determine if the packet should be accepted anyway. The fail filter in this example allows Device B’s interfaces to accept Dynamic Host Configuration Protocol (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.

 Copyright © 2018, Juniper Networks, Inc.

Figure 13 on page 255 shows the sample network.

Figure 13: Unicast RPF Sample Topoolgy

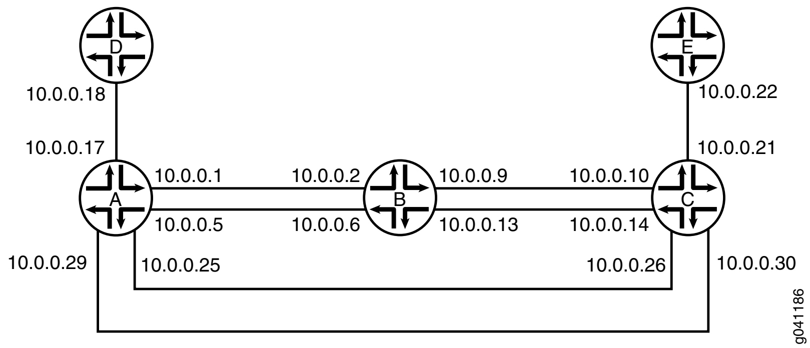

Configuration

CLI Quick Configuration

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

Device A

```
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces fe-0/0/2 unit 5 family inet address 10.0.0.5/30
set interfaces fe-0/0/1 unit 17 family inet address 10.0.0.17/30
set interfaces fe-0/1/1 unit 25 family inet address 10.0.0.25/30
set interfaces fe-1/1/1 unit 29 family inet address 10.0.0.29/30
set protocols ospf export send-direct
set protocols ospf area 0.0.0.0 interface fe-0/1/1.25
set protocols ospf area 0.0.0.0 interface fe-1/1/1.29
set policy-options policy-statement send-direct from protocol direct
set policy-options policy-statement send-direct from route-filter 10.0.0.16/30 exact
set policy-options policy-statement send-direct then accept
```

Device B

```
set interfaces fe-1/2/0 unit 2 family inet rpf-check fail-filter rpf-special-case-dhcp
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/1/1 unit 6 family inet rpf-check fail-filter rpf-special-case-dhcp
set interfaces fe-1/1/1 unit 6 family inet address 10.0.0.6/30
set interfaces fe-0/1/1 unit 9 family inet rpf-check fail-filter rpf-special-case-dhcp
set interfaces fe-0/1/1 unit 9 family inet address 10.0.0.9/30
set interfaces fe-0/1/0 unit 13 family inet rpf-check fail-filter rpf-special-case-dhcp
set interfaces fe-0/1/0 unit 13 family inet address 10.0.0.13/30
set protocols ospf area 0.0.0.0 interface fe-0/1/1.9
set protocols ospf area 0.0.0.0 interface fe-0/1/0.13
set routing-options forwarding-table unicast-reverse-path active-paths
set firewall filter rpf-special-case-dhcp term allow-dhcp from source-address 0.0.0.0/32
set firewall filter rpf-special-case-dhcp term allow-dhcp from destination-address
  255.255.255.255/32
set firewall filter rpf-special-case-dhcp term allow-dhcp then count rpf-dhcp-traffic
set firewall filter rpf-special-case-dhcp term allow-dhcp then accept
set firewall filter rpf-special-case-dhcp term default then log
set firewall filter rpf-special-case-dhcp term default then reject
```

Copyright © 2018, Juniper Networks, Inc.

Device C

```
set interfaces fe-1/2/0 unit 10 family inet address 10.0.0.10/30
set interfaces fe-0/0/2 unit 14 family inet address 10.0.0.14/30
set interfaces fe-1/0/2 unit 21 family inet address 10.0.0.21/30
set interfaces fe-1/2/2 unit 26 family inet address 10.0.0.26/30
set interfaces fe-1/2/1 unit 30 family inet address 10.0.0.30/30
set protocols ospf area 0.0.0.0 interface fe-1/2/0.10
set protocols ospf area 0.0.0.0 interface fe-0/0/2.14
set protocols ospf area 0.0.0.0 interface fe-1/2/2.26
set protocols ospf area 0.0.0.0 interface fe-1/2/1.30
```

Device D

```
set interfaces fe-1/2/0 unit 18 family inet address 10.0.0.18/30
```

Device E

```
set interfaces fe-1/2/0 unit 22 family inet address 10.0.0.22/30
```

Configuring Device A

Step-by-Step Procedure

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

To configure Device A:

1. Configure the interfaces.

   ```
   [edit interfaces]
   user@A# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

   user@A# set fe-0/0/2 unit 5 family inet address 10.0.0.5/30

   user@A# set fe-0/0/1 unit 17 family inet address 10.0.0.17/30

   user@A# set fe-0/1/1 unit 25 family inet address 10.0.0.25/30

   user@A# set fe-1/1/1 unit 29 family inet address 10.0.0.29/30
   ```

2. Configure OSPF.

   ```
   [edit protocols ospf]
   user@A# set export send-direct
   user@A# set area 0.0.0.0 interface fe-0/1/1.25
   user@A# set area 0.0.0.0 interface fe-1/1/1.29
   ```

3. Configure the routing policy.

   ```
   [edit policy-options policy-statement send-direct]
   user@A# set from protocol direct
   user@A# set from route-filter 10.0.0.16/30 exact
   user@A# set then accept
   ```

4. If you are done configuring Device A, commit the configuration.

   ```
   [edit]
   ```

 Copyright © 2018, Juniper Networks, Inc.

```
user@A# commit
```

Configuring Device B

Step-by-Step Procedure

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

To configure Device B:

1. Configure the interfaces.

   ```
   [edit interfaces]
   user@B# set fe-1/2/0 unit 2 family inet address 10.0.0.2/30

   user@B# set fe-1/1/1 unit 6 family inet address 10.0.0.6/30

   user@B# set fe-0/1/1 unit 9 family inet address 10.0.0.9/30

   user@B# set fe-0/1/0 unit 13 family inet address 10.0.0.13/30
   ```

2. Configure OSPF.

   ```
   [edit protocols ospf area 0.0.0.0]
   user@B# set interface fe-0/1/1.9
   user@B# set interface fe-0/1/0.13
   ```

3. Configure unicast RPF, and apply the optional fail filter.

   ```
   [edit interfaces]
   user@B# set fe-1/2/0 unit 2 family inet rpf-check fail-filter rpf-special-case-dhcp

   user@B# set fe-1/1/1 unit 6 family inet rpf-check fail-filter rpf-special-case-dhcp

   user@B# set fe-0/1/1 unit 9 family inet rpf-check fail-filter rpf-special-case-dhcp

   user@B# set fe-0/1/0 unit 13 family inet rpf-check fail-filter rpf-special-case-dhcp
   ```

4. (Optional) Configure the fail filter that gets evaluated if a packet fails the RPF check.

   ```
   [edit firewall filter rpf-special-case-dhcp]
   user@B# set term allow-dhcp from source-address 0.0.0.0/32
   user@B# set term allow-dhcp from destination-address 255.255.255.255/32
   user@B# set term allow-dhcp then count rpf-dhcp-traffic
   user@B# set term allow-dhcp then accept
   user@B# set term default then log
   user@B# set term default then reject
   ```

5. (Optional) Configure only active paths to be considered in the RPF check.

   This is the default behavior.

```
[edit routing-options forwarding-table]
user@B# set unicast-reverse-path active-paths
```

6. If you are done configuring Device B, commit the configuration.

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

Results

Confirm your configuration by issuing the **show firewall**, **show interfaces**, **show protocols**, **show routing-options**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Device A

```
user@A# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 10.0.0.1/30;
    }
  }
}
fe-0/0/2 {
  unit 5 {
    family inet {
      address 10.0.0.5/30;
    }
  }
}
fe-0/0/1 {
  unit 17 {
    family inet {
      address 10.0.0.17/30;
    }
  }
}
fe-0/1/1 {
  unit 25 {
    family inet {
      address 10.0.0.25/30;
    }
  }
}
fe-1/1/1 {
  unit 29 {
    family inet {
      address 10.0.0.29/30;
    }
  }
}

user@A# show protocols
```


Copyright © 2018, Juniper Networks, Inc.

```
ospf {
   export send-direct;
   area 0.0.0.0 {
      interface fe-0/1/1.25;
      interface fe-1/1/1.29;
   }
}
```

```
user@A# show policy-options
policy-statement send-direct {
   from {
      protocol direct;
      route-filter 10.0.0.16/30 exact;
   }
   then accept;
}
```

Device B

```
user@B# show firewall
filter rpf-special-case-dhcp {
   term allow-dhcp {
      from {
         source-address {
            0.0.0.0/32;
         }
         destination-address {
            255.255.255.255/32;
         }
      }
      then {
         count rpf-dhcp-traffic;
         accept;
      }
   }
   term default {
      then {
         log;
         reject;
      }
   }
}
user@B# show interfaces
fe-1/2/0 {
   unit 2 {
      family inet {
         rpf-check fail-filter rpf-special-case-dhcp;
         address 10.0.0.2/30;
      }
   }
}
fe-1/1/1 {
   unit 6 {
      family inet {
         rpf-check fail-filter rpf-special-case-dhcp;
         address 10.0.0.6/30;
      }
   }
```

```
}
fe-0/1/1 {
  unit 9 {
    family inet {
      rpf-check fail-filter rpf-special-case-dhcp;
      address 10.0.0.9/30;
    }
  }
}
fe-0/1/0 {
  unit 13 {
    family inet {
      rpf-check fail-filter rpf-special-case-dhcp;
      address 10.0.0.13/30;
    }
  }
}

user@B# show protocols
ospf {
  area 0.0.0.0 {
    interface fe-0/1/1.9;
    interface fe-0/1/0.13;
  }
}

user@B# show routing-options
forwarding-table {
  unicast-reverse-path active-paths;
}
```

Enter the configurations on Device C, Device D, and Device E, as shown in "CLI Quick Configuration" on page 255.

Verification

Confirm that the configuration is working properly.

- Confirm That Unicast RPF Is Enabled on page 260
- Confirm That the Source Addresses Are Blocked on page 261
- Confirm That the Source Addresses Are Unblocked on page 261

Confirm That Unicast RPF Is Enabled

Purpose Make sure that the interfaces on Device B have unicast RPF enabled.

 Copyright © 2018, Juniper Networks, Inc.

Action

```
user@B> show interfaces fe-0/1/0.13 extensive
 Logical interface fe-0/1/0.13 (Index 73) (SNMP ifIndex 553) (Generation 208)
    Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :             999390
     Output bytes  :            1230122
     Input  packets:              12563
     Output packets:              12613
    Local statistics:
     Input  bytes  :             998994
     Output bytes  :            1230122
     Input  packets:              12563
     Output packets:              12613
    Transit statistics:
     Input  bytes  :                396                   0 bps
     Output bytes  :                  0                   0 bps
     Input  packets:                  0                   0 pps
     Output packets:                  0                   0 pps
    Protocol inet, MTU: 1500, Generation: 289, Route table: 22
      Flags: Sendbcast-pkt-to-re, uRPF
      RPF Failures: Packets: 0, Bytes: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.0.0.12/30, Local: 10.0.0.13, Broadcast: 10.0.0.15,
Generation: 241
```

Meaning The **uRPF** flag confirms that unicast RPF is enabled on this interface.

Confirm That the Source Addresses Are Blocked

Purpose Use the **ping** command to make sure that Device B blocks traffic from unexpected source addresses.

Action From Device A, ping Device B's interfaces, using 10.0.0.17 as the source address.

```
user@A> ping 10.0.0.6 source 10.0.0.17
PING 10.0.0.6 (10.0.0.6): 56 data bytes
^C
--- 10.0.0.6 ping statistics ---
3 packets transmitted, 0 packets received, 100% packet loss
```

Meaning As expected, the ping operation fails.

Confirm That the Source Addresses Are Unblocked

Purpose Use the **ping** command to make sure that Device B does not block traffic when the RPF check is deactivated.

Action

1. Deactivate the RPF check on one of the interfaces.

Copyright © 2018, Juniper Networks, Inc.

2. Rerun the ping operation.

```
user@B> deactivate interfaces fe-1/1/1.6 family inet rpf-check

user@A> ping 10.0.0.6 source 10.0.0.17
PING 10.0.0.2 (10.0.0.2): 56 data bytes
64 bytes from 10.0.0.2: icmp_seq=0 ttl=63 time=1.316 ms
64 bytes from 10.0.0.2: icmp_seq=1 ttl=63 time=1.263 ms
^C
--- 10.0.0.2 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.263/1.289/1.316/0.027 ms
```

Meaning As expected, the ping operation succeeds.

See Also

- Understanding How Unicast Reverse Path Forwarding Prevents Spoofed IP Packet Forwarding on page 253

Related Documentation

- *Example: Enabling Indirect Next Hops on the Packet Forwarding Engine*

Enabling Source Class and Destination Class Usage

Source Class and Destination Class Usage

For interfaces that carry IPv4, IPv6, MPLS, or peer AS billing traffic, you can maintain packet counts based on the entry and exit points for traffic passing through your network. Entry and exit points are identified by source and destination prefixes grouped into disjoint sets defined as *source classes* and *destination classes*. You can define classes based on a variety of parameters, such as routing neighbors, autonomous systems, and route filters.

Source class usage (SCU) counts packets sent to customers by performing lookup on the IP source address. SCU makes it possible to track traffic originating from specific prefixes on the provider core and destined for specific prefixes on the customer edge. You must enable SCU accounting on both the inbound and outbound physical interfaces, and the route for the source of the packet must be in located in the forwarding table.

NOTE: SCU and DCU accounting do not work with directly connected interface routes. Source class usage does not count packets coming from sources with direct routes in the forwarding table because of software architecture limitations.

 Copyright © 2018, Juniper Networks, Inc.

Destination class usage (DCU) counts packets from customers by performing lookup of the IP destination address. DCU makes it possible to track traffic originating from the customer edge and destined for specific prefixes on the provider core router.

NOTE: We recommend that you stop the network traffic on an interface before you modify the DCU or SCU configuration for that interface. Modifying the DCU or SCU configuration without stopping the traffic might corrupt the DCU or SCU statistics. Before you restart the traffic after modifying the configuration, enter the **clear interfaces statistics** command.

Figure 14 on page 263 illustrates an Internet service provider (ISP) network. In this topology, you can use DCU to count packets customers send to specific prefixes. For example, you can have three counters, one per customer, that count the packets destined for prefix **210.210/16** and **220.220/16**.

You can use SCU to count packets the provider sends from specific prefixes. For example, you can count the packets sent from prefix **210.210/16** and **215.215/16** and transmitted on a specific output interface.

Figure 14: Prefix Accounting with Source and Destination Classes

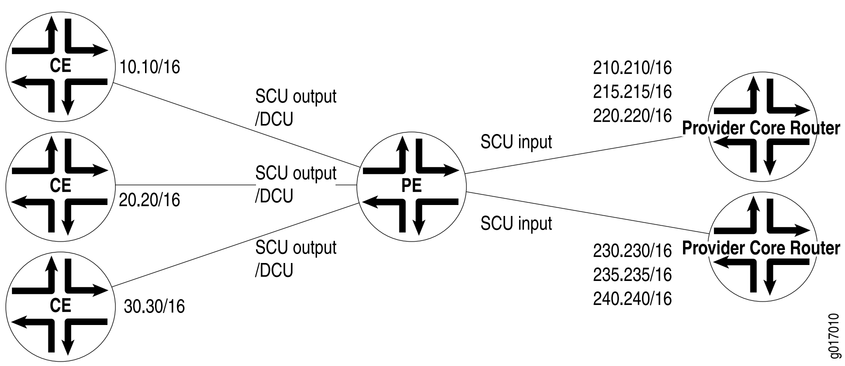

You can configure up to 126 source classes and 126 destination classes. For each interface on which you enable destination class usage and source class usage, the Junos OS maintains an interface-specific counter for each corresponding class up to the 126 class limit.

NOTE: 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.

Copyright © 2018, Juniper Networks, Inc.

NOTE:

Performing DCU accounting when an output service is enabled produces inconsistent behavior in the following configuration:

- Both SCU input and DCU are configured on the packet input interface.
- SCU output is configured on the packet output interface.
- Interface services is enabled on the output interface.

For an incoming packet with source and destination prefixes matching the SCU and DCU classes respectively configured in the router, both SCU and DCU counters will be incremented. This behavior is not harmful or negative. However, it is inconsistent with non-serviced packets, in that only the SCU count will be incremented (because the SCU class ID will override the DCU class ID in this case).

To enable packet counting on an interface, include the **accounting** statement:

```
accounting {
  destination-class-usage;
  source-class-usage {
    direction;
  }
}
```

direction can be one of the following:

- **input**—Configure at least one expected ingress point.
- **output**—Configure at least one expected egress point.
- **input output**—On a single interface, configure at least one expected ingress point and one expected egress point.

You can include these statements at the following hierarchy levels:

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

For SCU to work, you must configure at least one input interface and at least one output interface.

The ability to count a single packet for both SCU and DCU accounting depends on the underlying physical interface.

- For traffic over MPC/MIC interfaces , a single incoming packet is counted for both SCU and DCU accounting if both SCU and DCU are configured. To ensure the outgoing packet is counted, include the **source-class-usage output** statements in the configuration of the outgoing interface.

 Copyright © 2018, Juniper Networks, Inc.

- For traffic over DPC interfaces, an incoming packet is counted only once, and SCU takes priority over DCU. This means that when a packet arrives on an interface on which you include the **source-class-usage input** and **destination-class-usage** statements in the configuration, and when the source and destination both match accounting prefixes, the Junos OS associates the packet with the source class only.

For traffic over MPC interfaces , SCU and DCU accounting is performed after output filters are evaluated. If a packet matches a firewall filter match condition, the packet is included in SCU or DCU accounting except in the case where the action of the matched term is **discard**.

On T Series, M120, and M320 routers, the source class and destination classes are not carried across the router fabric. The implications of this are as follows:

- On T Series, M120, and M320 routers, SCU and DCU accounting is performed before the packet enters the fabric.
- On M7i, M10i, M120, and M320 routers, on MX Series routers with non-MPC, and on T Series routers, SCU and DCU accounting is performed before output filters are evaluated. Consequently, if a packet matches a firewall filter match condition, the packet is included in SCU or DCU accounting; the packet is counted for any term action (including the **discard** action).
- On M120, M320, and T Series routers, the **destination-class** and **source-class** statements are supported at the **[edit firewall family *family-name* filter *filter-name* term *term-name* from]** hierarchy level only for the filter applied to the forwarding table. On M7i, M10i, and MX Series routers, these statements are supported.

Once you enable accounting on an interface, the Junos OS maintains packet counters for that interface, with separate counters for **inet**, **inet6**, and **mpls** protocol families. You must then configure the source class and destination class attributes in policy action statements, which must be included in forwarding-table export policies.

NOTE: When configuring policy action statements, you can configure only one source class for each matching route. In other words, more than one source class cannot be applied to the same route.

In Junos OS Release 9.3 and later, you can configure SCU accounting for Layer 3 VPNs configured with the **vrf-table-label** statement. Include the **source-class-usage** statement at the **[edit routing-instances *routing-instance-name* vrf-table-label]** hierarchy level. The **source-class-usage** statement at this hierarchy level is supported only for the virtual routing and forwarding (VRF) instance type.

NOTE: DCU counters cannot be enabled on the label-switched interface (LSI) that is created dynamically when the **vrf-table-label** statement is configured within a VRF. For more information, see the *Junos OS VPNs Library for Routing Devices*.

For a complete discussion about source and destination class accounting profiles, see the *Network Management and Monitoring Guide*. For more information about MPLS, see the *MPLS Applications Feature Guide*.

Enabling Source Class and Destination Class Usage

Figure 15: Prefix Accounting with Source and Destination Classes

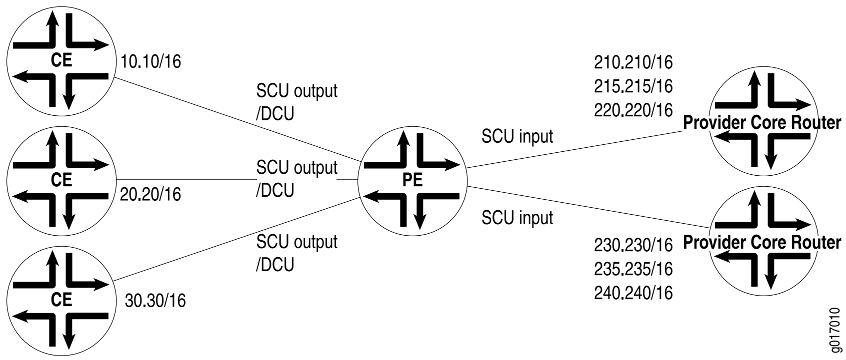

Configure DCU and SCU output on one interface:

```
[edit]
interfaces {
   so-6/1/0 {
      unit 0 {
         family inet {
            accounting {
               destination-class-usage;
               source-class-usage {
                  output;
               }
            }
         }
      }
   }
}
```

1. **Complete SCU Configuration**

   Source routers A and B use loopback addresses as the prefixes to be monitored. Most of the configuration tasks and actual monitoring occur on transit Router SCU.

   The loopback address on Router A contains the origin of the prefix that is to be assigned to source class A on Router SCU. However, no SCU processing happens on this router. Therefore, configure Router A for basic OSPF routing and include your loopback interface and interface **so-0/0/2** in the OSPF process.

2. **Router A**

```
[edit]
interfaces {
   so-0/0/2 {
      unit 0 {
         family inet {
            address 10.255.50.2/24;
```

 Copyright © 2018, Juniper Networks, Inc.

```
            }
         }
      }
      lo0 {
         unit 0 {
            family inet {
               address 10.255.192.10/32;
            }
         }
      }
   }
   protocols {
      ospf {
         area 0.0.0.0 {
            interface so-0/0/2.0;
            interface lo0.0;
         }
      }
   }
```

3. **Router SCU**

   Last, apply the policy to the forwarding table.

   Router SCU handles the bulk of the activity in this example. On Router SCU, enable source class usage on the inbound and outbound interfaces at the **[edit interfaces *interface-name* unit *unit-number* family inet accounting]** hierarchy level. Make sure you specify the expected traffic: input, output, or, in this case, both.

   Next, configure a route filter policy statement that matches the prefixes of the loopback addresses from routers A and B. Include statements in the policy that classify packets from Router A in one group named **scu-class-a** and packets from Router B in a second class named **scu-class-b**. Notice the efficient use of a single policy containing multiple terms.

```
[edit]
interfaces {
   so-0/0/1 {
      unit 0 {
         family inet {
            accounting {
               source-class-usage {
                  input;
                  output;
               }
            }
            address 10.255.50.1/24;
         }
      }
   }
   so-0/0/3 {
      unit 0 {
         family inet {
            accounting {
               source-class-usage {
                  input;
```

```
                output;
              }
            }
            address 10.255.10.3/24;
          }
        }
      }
      lo0 {
        unit 0 {
          family inet {
            address 10.255.6.111/32;
          }
        }
      }
    }
    protocols {
      ospf {
        area 0.0.0.0 {
          interface so-0/0/1.0;
          interface so-0/0/3.0;
        }
      }
    }
    routing-options {
      forwarding-table {
        export scu-policy;
      }
    }
    policy-options {
      policy-statement scu-policy {
        term 0 {
          from {
            route-filter 10.255.192.0/24 orlonger;
          }
          then source-class scu-class-a;
        }
        term 1 {
          from {
            route-filter 10.255.165.0/24 orlonger;
          }
          then source-class scu-class-b;
        }
      }
    }
```

4. **Router B**

   Just as Router A provides a source prefix, Router B's loopback address matches the prefix assigned to **scu-class-b** on Router SCU. Again, no SCU processing happens on this router, so configure Router B for basic OSPF routing and include your loopback interface and interface **so-0/0/4** in the OSPF process.

```
    interfaces {
      so-0/0/4 {
        unit 0 {
          family inet {
```

 Copyright © 2018, Juniper Networks, Inc.

```
                address 10.255.10.4/24;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 10.255.165.226/32;
            }
        }
    }
}
protocols {
    ospf {
        area 0.0.0.0 {
            interface so-0/0/4.0;
            interface lo0.0;
        }
    }
}
```

5. **Enabling Packet Counting for Layer 3 VPNs**

   You can use SCU and DCU to count packets on Layer 3 VPNs. To enable packet counting for Layer 3 VPN implementations at the egress point of the MPLS tunnel, you must configure a virtual loopback tunnel interface (**vt**) on the PE router, map the virtual routing and forwarding (VRF) instance type to the virtual loopback tunnel interface, and send the traffic received from the VPN out the source class output interface, as shown in the following example:

   Configure a virtual loopback tunnel interface on a provider edge router equipped with a tunnel PIC:

```
[edit interfaces]
vt-0/3/0 {
    unit 0 {
        family inet {
            accounting {
                source-class-usage {
                    input;
                }
            }
        }
    }
}
```

6. Map the VRF instance type to the virtual loopback tunnel interface.

   In Junos OS Release 9.3 and later, you can configure SCU accounting for Layer 3 VPNs configured with the **vrf-table-label** statement. Include the **source-class-usage** statement at the **[edit routing-instances *routing-instance-name* vrf-table-label]** hierarchy level. The **source-class-usage** statement at this hierarchy level is supported only for the virtual routing and forwarding (VRF) instance type. DCU is not supported when the **vrf-table-label** statement is configured. For more information, see the *Junos OS VPNs Library for Routing Devices*.

   ```
   [edit routing-instances]
   VPN-A {
     instance-type vrf;
     interface at-2/1/1.0;
     interface vt-0/3/0.0;
     route-distinguisher 10.255.14.225:100;
     vrf-import import-policy-A;
     vrf-export export-policy-A;
     protocols {
       bgp {
         group to-r4 {
           local-address 10.27.253.1;
           peer-as 400;
           neighbor 10.27.253.2;
         }
       }
     }
   }
   ```

7. Send traffic received from the VPN out the source class output interface:

   ```
   [edit interfaces]
   at-2/1/0 {
     unit 0 {
       family inet {
         accounting {
           source-class-usage {
             output;
           }
         }
       }
     }
   }
   ```

   For more information about VPNs, see the *Junos OS VPNs Library for Routing Devices*. For more information about virtual loopback tunnel interfaces, see the *Junos OS Services Interfaces Library for Routing Devices*.

See Also

- accounting on page 413
- *destination-classes*
- family on page 607
- forward-and-send-to-re on page 622

 Copyright © 2018, Juniper Networks, Inc.

- *source-classes*
- targeted-broadcast on page 1010
- unit on page 1056

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 forwarded to the Routing Engine.

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: Targeted broadcast does not work when the targeted broadcast option **forward-and-send-to-re** and the traffic sampling option **sampling** are configured on the same egress interface of an M320 router, a T640 router, or an MX960 router. To overcome this scenario, you must either disable one of the these options or enable the **sampling** option with the targeted broadcast option **forward-only** on the egress interface. For information about traffic sampling, see *Configuring Traffic Sampling*.

Copyright © 2018, Juniper Networks, Inc.

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 272
- **targeted-broadcast on page 1010**

Configuring Targeted Broadcast

The following sections explain how to configure targeted broadcast on an egress interface and its options:

Configuring Targeted Broadcast and Its Options

You can configure targeted broadcast on an egress interface with different options. You can either allow the IP packets destined for a Layer 3 broadcast address to be forwarded on the egress interface and to send a copy of the IP packets to the Routing Engine or you can allow the IP packets to be forwarded on the egress interface only. Note that the packets are broadcast only if the egress interface is a LAN interface.

To configure targeted broadcast and its options:

1. Configure the physical interface.

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

2. Configure the logical unit number at the **[edit interfaces *interface-name*** hierarchy level.

   [edit interfaces *interface-name*]
   user@host# **set unit *logical-unit-number***

3. Configure the protocol family as inet at the **[edit interfaces *interface-name* unit *interface-unit-number*** hierarchy level.

   [edit interfaces *interface-name* unit *interface--unit-number*]
   user@host# **set family inet**

4. Configure targeted broadcast at the **[edit interfaces *interface-name* unit *interface-unit-number* family inet** hierarchy level

   [edit interfaces *interface-name* unit *interface--unit-number* family inet]

Copyright © 2018, Juniper Networks, Inc.

user@host# **set targeted-broadcast**

5. Specify one of the following options as per requirement:

   - To allow IP packets destined for a Layer 3 broadcast address to be forwarded on the egress interface and to send a copy of the IP packets to the Routing Engine.

     [edit interfaces *interface-name* unit *interface-unit-number* family inet
       targeted-broadcast]
     user@host# **set forward-and-send-to-re**

   - To allow IP packets to be forwarded on the egress interface only.

     [edit interfaces *interface-name* unit *interface-unit-number* family inet
       targeted-broadcast]
     user@host# **set forward-only**

NOTE: Targeted broadcast does not work when the targeted broadcast option **forward-and-send-to-re** and the traffic sampling option **sampling** are configured on the same egress interface of an M320 router, a T640 router, or an MX960 router. To overcome this scenario, you must either disable one of the these options or enable the **sampling** option with the targeted broadcast option **forward-only** on the egress interface. For information about traffic sampling, see *Configuring Traffic Sampling*.

Display Targeted Broadcast Configuration Options

The following topics display targeted broadcast configuration with its various options:

Forward IP Packets On the Egress Interface and To the Routing Engine

Purpose Display the configuration when targeted broadcast is configured on the egress interface to forward the IP packets on the egress interface and to send a copy of the IP packets to the Routing Engine.

Action To display the configuration run the **show** command at the **[edit interfaces *interface-name* unit *interface-unit-number* family inet]** where the interface name is ge-2/0/0, the unit value is set to 0, the protocol family is set to inet.

[edit interfaces *interface-name* unit *interface-unit-number* family inet]
user@host#**show**

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

Copyright © 2018, Juniper Networks, Inc.

Forward IP Packets On the Egress Interface Only

Purpose Display the configuration when targeted broadcast is configured on the egress interface to forward the IP packets on the egress interface only.

Action To display the configuration run the **show** command at the **[edit interfaces *interface-name* unit *interface-unit-number* family inet]** where the interface name is ge-2/0/0, the unit value is set to 0, the protocol family is set to inet.

```
[edit interfaces interface-name unit interface-unit-number family inet]
user@host#show
targeted-broadcast {
  forward-only;
}
```

Related Documentation

- **targeted-broadcast on page 1010**
- Understanding Targeted Broadcast on page 271

 Copyright © 2018, Juniper Networks, Inc.

CHAPTER 5

Configuring Circuit and Translational Cross-Connects

Circuit and Translational Cross-Connects Overview

Circuit cross-connect (CCC) and translational cross-connect (TCC) allow you to configure transparent connections between two circuits, where a circuit can be a Frame Relay data-link connection identifier (DLCI), an Asynchronous Transfer Mode (ATM) virtual circuit (VC), a Point-to-Point Protocol (PPP) interface, a Cisco High-level Data Link Control (HDLC) interface, or a Multiprotocol Label Switching (MPLS) label-switched path (LSP).

Using CCC or TCC, packets from the source circuit are delivered to the destination circuit with, at most, the Layer 2 address being changed. No other processing, such as header checksums, time-to-live (TTL) decrementing, or protocol processing, is done.

To connect interfaces of the same type, use CCC. To connect unlike interfaces, use TCC.

Copyright © 2018, Juniper Networks, Inc.

CCC and TCC circuits fall into three categories: logical interfaces, which include ATM VCs and Frame Relay DLCIs; physical interfaces, which include PPP and Cisco HDLC; and paths, which include LSPs. The three circuit categories provide three types of cross-connect:

- Layer 2 switching (interface-to-interface)—Cross-connects between logical interfaces provide what is essentially Layer 2 switching.
- MPLS tunneling (interface-to-LSP)—Cross-connects between interfaces and LSPs allow you to connect two distant interface circuits by creating MPLS tunnels that use LSPs as the conduit.
- LSP stitching (LSP-to-LSP)—Cross-connects between LSPs provide a way to “stitch” together two label-switched paths, including paths that fall in two different traffic engineering database (TED) areas.

The cross-connect is bidirectional, so packets received on the first interface are transmitted out the second interface, and those received on the second interface are transmitted out the first interface.

For most CCC connections that connect interfaces, the interfaces must be of the same type; that is, ATM to ATM, Frame Relay to Frame Relay, PPP to PPP, or Cisco HDLC to Cisco HDLC.

ATM-to-Ethernet interworking cross-connect circuits connect logical interfaces configured on an ATM2 and Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet IQ2 and IQ2-E physical interfaces.

For all TCC connections that connect interfaces, the interfaces can be of unlike types. Mainly, TCC is used for Layer 2.5 virtual private networks (VPNs), but it can also be used as a simple “unlike circuit” switch.

Switching cross-connects join logical interfaces to form what is essentially Layer 2 switching.

Figure 16 on page 276 illustrates a Layer 2 switching circuit cross-connect. In this topology, Router A and Router C have Frame Relay connections to Router B, which is a Juniper Networks router. CCC allows you to configure Router B to act as a Frame Relay (Layer 2) switch. To do this, configure a circuit from Router A to Router C that passes through Router B, effectively configuring Router B as a Frame Relay switch with respect to these routers. This configuration allows Router B to transparently switch packets (frames) between Router A and Router C without regard to the packets’ contents or the Layer 3 protocols. The only processing that Router B performs is to translate DLCI 600 to 750.

Figure 16: Layer 2 Switching Circuit Cross-Connect

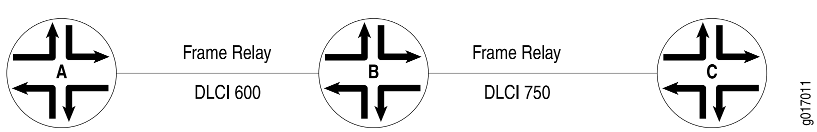

If the Router A–to–Router B and Router B–to–Router C circuits are PPP, for example, the Link Control Protocol and Network Control Protocol exchanges occur between Router A and Router C. These messages are handled transparently by Router B, allowing

 Copyright © 2018, Juniper Networks, Inc.

Router A and Router C to use various PPP options (such as header or address compression and authentication) that Router B might not support. Similarly, Router A and Router C exchange keepalives, providing circuit-to-circuit connectivity status.

You can configure Layer 2 switching cross-connects on PPP, Cisco HDLC, Frame Relay, Ethernet CCC, Ethernet VLAN, and ATM circuits. With CCC, only like interfaces can be connected in a single cross-connect. With TCC, unlike interfaces can be connected in a single cross-connect. In Layer 2 switching cross-connects, the exchanges take place between point-to-point links.

This chapter discusses the Layer 2 switching cross-connect configuration tasks. For information about MPLS tunneling and LSP stitching, see the *Junos OS MPLS Applications Library for Routing Devices*.

For information about Layer 2 and Layer 2.5 VPNs, see the *Junos OS VPNs Library for Routing Devices*.

For restrictions for MPLS on QFX switches, see *MPLS Feature Support on QFX Series and EX4600 Switches*.

Related Documentation

- Defining the Encapsulation for Switching Cross-Connects on page 277
- Defining the Connection for Switching Cross-Connects on page 281
- Configuring MPLS for Switching Cross-Connects on page 281
- Configuring IS-IS or MPLS Traffic for TCC Interfaces on page 281
- Configuring ATM-to-Ethernet Interworking on page 282
- Example: Configuring a CCC over Frame Relay Encapsulated Interface on page 287
- Example: Configuring a TCC on page 288
- Example: Configuring CCC over Aggregated Ethernet on page 290
- Example: Configuring a Remote LSP CCC over Aggregated Ethernet on page 291
- Example: Configuring ATM-to-Ethernet Interworking on page 294

Defining the Encapsulation for Switching Cross-Connects

Defining the Encapsulation for Switching Cross-Connects

To configure Layer 2 or Layer 2.5 switching cross-connects, configure the CCC or TCC encapsulation on the router that is acting as the switch (Router B in Figure 16 on page 276).

NOTE: When you use CCC encapsulation, you can configure the ccc family only. Likewise, when you use TCC encapsulation, you can configure the tcc family only.

This section contains the following topics:

Configuring PPP or Cisco HDLC Circuits

For PPP or Cisco HDLC circuits, specify the encapsulation by including the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level. This statement configures the entire physical device. For these circuits to work, you must configure a logical interface unit 0.

```
[edit interfaces interface-name]
encapsulation (ppp-ccc | cisco-hdlc-ccc | ppp-tcc | cisco-hdlc-tcc);
unit 0;
```

Configuring ATM Circuits

For ATM circuits, include the **vpi** statement **[edit interfaces *interface-name* atm-options]** hierarchy level:

```
[edit interfaces at-fpc/pic/port]
atm-options {
  vpi vpi-identifier;
}
```

On the logical interface, include the following statements:

```
point-to-point;
encapsulation (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-tcc-vc-mux | atm-tcc-snap);
vci vpi-identifier.vci-identifier;
```

You can include the logical interface statements at the following hierarchy levels:

- **[edit interfaces at-*fpc/pic/port* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces at-*fpc/pic/port* unit *logical-unit-number*]**

For each VC, configure whether it is a circuit or a regular logical interface. The default interface type is point-to-point.

Configuring Frame Relay Circuits

For Frame Relay circuits, include the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level:

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

On the logical interface, include the following statements:

```
point-to-point;
```


Copyright © 2018, Juniper Networks, Inc.

```
encapsulation type;
dlci dlci-identifier;
```

You can include the logical interface statements at the following hierarchy levels:

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

The encapsulation type can be one of the following:

- Flexible Frame Relay (**flexible-frame-relay**)—Intelligent queuing (IQ) interfaces can use flexible Frame Relay encapsulation. You use flexible Frame Relay encapsulation when you want to configure multiple per-unit Frame Relay encapsulations. This encapsulation type allows you to configure any combination of TCC, CCC, and standard Frame Relay encapsulations on a single physical port. Also, each logical interface can have any DLCI value from 1 through 1022.
- Frame Relay CCC version (**frame-relay-ccc**)—For E1, E3, SONET/SDH, T1, and T3 interfaces, this encapsulation type is the same as standard Frame Relay for DLCIs 0 through 511. DLCIs 512 through 1022 are dedicated to CCC. The logical interface must also have **frame-relay-ccc** encapsulation. When you use this encapsulation type, you can configure the **ccc** family only.
- Frame Relay TCC version (**frame-relay-tcc**)—Similar to Frame Relay CCC and has the same configuration restrictions, but used for circuits with different media on either side of the connection.
- Extended CCC version (**extended-frame-relay-ccc**)—This encapsulation type allows you to dedicate DLCIs 1 through 1022 to CCC. The logical interface must have **frame-relay-ccc** encapsulation. When you use this encapsulation type, you can configure the **ccc** family only.
- Extended TCC version (**extended-frame-relay-tcc**)—Similar to extended Frame Relay CCC, this encapsulation type allows you to dedicate DLCIs 1 through 1022 to TCC, which is used for circuits with different media on either side of the connection.
- Port CCC version (**frame-relay-port-ccc**)—Defined in the IETF document *Frame Relay Encapsulation over Pseudo-Wires* (expired December 2002). This encapsulation type allows you to transparently carry all the DLCIs between two customer edge (CE) routers without explicitly configuring each DLCI on the two provider edge (PE) routers with Frame Relay transport. The connection between the two CE routers can be either user-to-network interface (UNI) or network-to-network interface (NNI); this is completely transparent to the PE routers. The logical interface does not require an encapsulation statement. When you use this encapsulation type, you can configure the **ccc** family only.

For each DLCI, configure whether it is a circuit or a regular logical interface. The DLCI for regular interfaces must be from 1 through 511. For CCC and TCC interfaces, it must be from 512 through 1022. This restriction does not apply to IQ interfaces. The default interface type is point to point.

Copyright © 2018, Juniper Networks, Inc.

Configuring Ethernet CCC Circuits

You can configure Ethernet CCC encapsulation on Fast Ethernet, Gigabit Ethernet, and aggregated Ethernet interfaces.

NOTE: CCC over aggregated Ethernet requires an M Series Enhanced Flexible PIC Concentrator (FPC).

For Ethernet CCC circuits, specify the encapsulation by including the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level. This statement configures the entire physical device.

```
[edit interfaces interface-name]
encapsulation ethernet-ccc;
unit logical-unit-number {
    ...
}
[edit interfaces aex]
encapsulation ethernet-ccc;
unit logical-unit-number {
    ...
}
```

Configuring Ethernet VLAN Circuits

You can configure Ethernet virtual local area network (VLAN) circuits on Fast Ethernet, Gigabit Ethernet, and aggregated Ethernet interfaces. For Ethernet VLAN circuits, specify the encapsulation by including the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level. This statement configures the entire physical device. You must also enable VLAN tagging. To do this, include the following statements:

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

On the logical interface, include the following statements:

```
encapsulation vlan-ccc;
vlan-id number;
```

You can include the logical interface 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*]**

Ethernet interfaces in VLAN mode can have multiple logical interfaces. For encapsulation type **vlan-ccc**, VLAN IDs 1 through 511 are reserved for normal VLANs, and VLAN IDs

Copyright © 2018, Juniper Networks, Inc.

512 through 1023 are reserved for CCC VLANs. For encapsulation type **extended-vlan-ccc**, VLAN IDs 1 through 4094 are valid. VLAN ID 0 is reserved for tagging the priority of frames.

See Also

- Figure 16 on page 276
- Circuit and Translational Cross-Connects Overview on page 275

Defining the Connection for Switching Cross-Connects

To configure Layer 2 switching cross-connects, define the connection between the two circuits. You configure this on the router that is acting as the switch (Router B in Figure 16 on page 276). The connection joins the interface that comes from the circuit's source to the interface that leads to the circuit's destination. When you specify the interface names, include the logical portion of the name, which corresponds to the logical unit number. The cross-connect is bidirectional, so packets received on the first interface are transmitted out the second interface, and those received on the second interface are transmitted out the first interface.

```
[edit protocols]
connections {
  remote-interface-switch connection-name {
    interface interface-name.unit-number;
  }
  lsp-switch connection-name {
    transmit-lsp lsp-number;
    receive-lsp lsp-number;
  }
}
```

Configuring MPLS for Switching Cross-Connects

For Layer 2 switching cross-connects to work, you must configure MPLS. The following is a minimal MPLS configuration:

```
[edit protocols]
mpls {
  interface (interface-name | all);
}
```

For more information, see the *MPLS Applications Feature Guide*.

Related Documentation

- Defining the Connection for Switching Cross-Connects on page 281

Configuring IS-IS or MPLS Traffic for TCC Interfaces

Layer 2.5 VPNs on T Series, M120, MX Series, and M320 routers support IPv4, IS-IS, and MPLS traffic types. By default, IPv4 traffic runs on T Series, M120, MX Series, and M320 routers and over TCC interfaces. To configure IS-IS (ISO traffic) or MPLS traffic on Layer 2.5 VPNs, you must configure the same traffic type on both ends of the Layer 2.5 VPN.

Copyright © 2018, Juniper Networks, Inc.

NOTE: Some platform and FPC combinations can not pass TCC encapsulated ISO traffic. See *Platforms/FPCs That Cannot Forward TCC Encapsulated ISO Traffic* for details.

To specify which traffic can run over a TCC interface, include the **protocols** statement with the appropriate value (**inet**, **mpls**, and **iso**) at the **[edit interfaces *interface-name* unit *logical-unit-number* family tcc]** hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family tcc]
protocols [ inet iso mpls ];
```

NOTE: Layer 2.5 VPNs running on M Series Multiservice Edge Routers support only IPv4 traffic. IPv6 is not supported on Layer 2.5 VPNs.

When enabling ISO over a Layer 2.5 VPN that is configured on a CE Ethernet interface, you must also include the **point-to-point** statement at the **[edit protocols isis interface *interface-name*]** hierarchy level:

```
[edit protocols isis interface interface-name]
point-to-point;
```

For more information about Layer 2.5 VPNs, see the *Junos OS VPNs Library for Routing Devices* and the *Translational Cross-Connect and Layer 2.5 VPNs Feature Guide*.

Configuring ATM-to-Ethernet Interworking

ATM-to-Ethernet Interworking

The ATM-to-Ethernet interworking feature is useful where ATM2 interfaces are used to terminate ATM DSLAM traffic. The ATM traffic can be forwarded with encapsulation type **ccc** (circuit cross-connect) to a local or remote Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet IQ2 and IQ2-E interface or label-switched path (LSP). The ATM VPI and VCI are converted to stacked VLAN inner and outer VLAN tags.

Copyright © 2018, Juniper Networks, Inc.

These ATM-to-Ethernet interworking circuits can be mapped to individual logical interfaces configured on an ATM2 IQ interface and Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet IQ2 and IQ2-E physical interface.

The ATM-to-Ethernet interworking cross-connect essentially provides Layer 2 switching, and statistics are reported at the logical interface level.

During conversion from ATM to Ethernet, the least significant 12 bits of the ATM cell VCI are copied to the Ethernet frame inner VLAN tag. Cells received on an ATM logical interface configured with encapsulation type **vlan-vci-ccc** and falling within the configured VCI range are reassembled into packets and forwarded to a designated Ethernet logical interface that is configured with encapsulation type **vlan-vci-ccc**.

During conversion from Ethernet to ATM, the Ethernet frame inner VLAN tags that fall within the configured range, are copied to the least significant 12 bits of the ATM cell VCI. The ATM logical interface uses its configured VPI when segmenting the Ethernet packets into cells.

ATM-to-Ethernet interworking is supported on M120, M320, and T Series routers.

ATM-to-Ethernet interworking is supported on MX Series routers with aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces. This feature is available on all Enhanced Queuing (EQ) DPCs and Enhanced DPCS for MX Series routers.

NOTE: This feature is *not* supported on MX Series routers with ATM interfaces.

For more information on MX Series ATM-to-Ethernet interworking, see the *MX Series Solutions Guide*.

The following sections discuss ATM-to-Ethernet interworking:

Enabling ATM-to-Ethernet Interworking

To enable the ATM-to-Ethernet interworking cross-connect function, include the **vlan-vci-tagging** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
vlan-vci-tagging;
```

Configuring the Ethernet Interface

Configure the Ethernet or aggregated Ethernet physical interface by including the **encapsulation** statement with the **vlan-vci-ccc** option at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
encapsulation vlan-vci-ccc;
```

When the encapsulation type **vlan-vci-ccc** is configured on the physical interface, all logical interfaces configured on the Ethernet interface must also have the encapsulation type set to **vlan-vci-ccc**.

Copyright © 2018, Juniper Networks, Inc.

Configuring Ethernet Encapsulation

Configure the Ethernet logical interface by including the **encapsulation** statement with the **vlan-vci-ccc** option at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number]
encapsulation vlan-vci-ccc;
```

The chassis configuration cannot contain the **atm-l2circuit-mode** statement if any logical interfaces are configured with the **vlan-vci-ccc** encapsulation option.

Configuring the Outer VLAN Identifier

Configure the Ethernet logical interface outer VLAN ID by including the **vlan-id** statement specifying the outer VLAN ID at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number]
vlan-id outer-vlan-identifier;
```

It is the administrator's responsibility to ensure that the outer VLAN tag and VPI match and the inner VLAN tags fall within the VCI range of the VPI.

The allowable VPI range is from 0 to 255. So the outer VLAN tags must not be configured for values above 255.

Configuring the Inner VLAN Identifier Range

Configure the Ethernet logical interface inner VLAN ID range by including the **inner-vlan-id-range** statement and specifying the starting VLAN ID and ending VLAN ID at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number]
inner-vlan-id-range start start-id end end-id;
```

VLAN IDs 0 and 4095 are reserved by IEEE 801.1q and must not be used for the inner or outer VLAN ID.

VCIs 0 through 31 are reserved for ATM management purposes by convention. Therefore inner VLAN IDs 1 through 31 should not be used.

VLAN ID 1 might be used by Ethernet switches for certain bridge management services, so using VLAN ID 1 for the inner or outer VLAN ID is discouraged.

Configuring the Physical Interface VPI

Configure the ATM physical interface VPI by including the **vpi** statement at the **[edit interfaces *interface-name* atm-options]** hierarchy level:

```
[edit interfaces interface-name atm-options]
vpi virtual-path-identifier;
```

VPI 0 is reserved, and must not be used.

 Copyright © 2018, Juniper Networks, Inc.

ATM F4/F5 OAM is not supported for VPIs used in ATM-to-Ethernet interworking cross-connects. Any F4/F5 OAM cells received are discarded.

Only one logical interface may be declared per virtual path specified in the **atm-options** statement hierarchy.

It is not necessary to dedicate all the VPIs of an ATM2 interface for ATM-to-Ethernet interworking cross-connects.

Configuring the ATM Logical Interface

Configure the ATM logical interface by including the **encapsulation** statement and specifying the encapsulation type **vlan-vci-ccc** at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number]
encapsulation vlan-vci-ccc;
```

An ATM logical interface configured with the encapsulation type **vlan-vci-ccc** only supports the **epd-threshold**, **shaping**, **traps | no-traps**, **disable**, and **description** statements. No other configuration statements are supported. ATM interface CoS features are not supported by logical interfaces configured with the encapsulation type **vlan-vci-ccc**.

The ATM2 OC48 PIC does not support the encapsulation type **vlan-vci-ccc**.

The encapsulation type **vlan-vci-ccc** only supports the **ccc** protocol family. Attempts to configure any other interface protocol family are rejected.

Configuring the Protocol Family

Configure the ATM logical interface protocol family by including the **family** statement and specifying the **ccc** option at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level:

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

Configuring the Logical Interface VPI

Configure the ATM logical interface virtual path identifier by including the **vpi** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number]
vpi virtual-path-identifier;
```

VPI 0 is reserved, and must not be used.

It is the administrator’s responsibility to ensure the outer VLAN tag and VPI match and the inner VLAN tags fall within the VCI range of the VPI.

Once a VPI is used in an ATM-to-Ethernet interworking cross-connect, it cannot be used with any other logical interface, even if the ***vpi.vci*** value falls outside the VCI range for the cross-connect.

Copyright © 2018, Juniper Networks, Inc.

Configuring the Logical Interface VCI

Configure the ATM logical interface virtual channel identifier range by including the **vci-range** statement and specifying the starting VCI and ending VCI at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level:

[edit interfaces *interface-name* unit *logical-unit-number*]
vci-range start *start-vci* end *end-vci*;

Do not use VCIs 0 through 31, which are reserved for ATM management purposes by convention.

See Also

- encapsulation on page 573
- family on page 607
- inner-vlan-id-range on page 673
- vci-range on page 1074
- vlan-id on page 1076
- vlan-vci-tagging on page 1095
- vpi (ATM CCC Cell-Relay Promiscuous Mode) on page 1095
- vpi (Logical Interface and Interworking) on page 1097

ATM-To-Ethernet Interworking on ATM MICs

ATM-to-Ethernet interworking supports transmission of ATM packets over Ethernet. It specifically provides support for exchange of Layer 2 and Layer 3 Protocol Data Units (PDUs) between ATM and Ethernet domains. On MX Series 3D Universal Edge Routers with ATM MICs, you can exchange Ethernet frames between ATM and Ethernet domains over a MPLS pseudowire or a Layer 2 cross-connect by using translational cross connect (TCC). For more information about TCC, see "Circuit and Translational Cross-Connects Overview" on page 275.

Consider the following basic ATM-to-Ethernet Interworking topology where the provider edge router PE1 is connected to an ATM domain and the Provider Edge router PE2 is connected to an Ethernet domain (see Figure 1). The customer edge routers CE1 and CE2 are customer-managed devices. The PE routers are connected by means of an MPLS pseudowire. The ATM traffic on the PE1–CE1 link can comprise untagged Ethernet frames over ATM format. The Ethernet traffic on PE2–CE2 link can comprise untagged, single-VLAN or double-VLAN tagged Ethernet frames depending on the configuration of the PE2 router.

For ATM-to-Ethernet Interworking, the virtual path identifier (VPI) and virtual circuit identifier (VCI) values on the ATM link are mapped to the outer and inner VLAN tag on the Ethernet link. Mapping implies that either the same value is copied or a one-to-one translation is performed. If VLAN translation is enabled, instead of copying the value, a one-to-one translation is performed on the Ethernet facing PE, using a lookup table.

 Copyright © 2018, Juniper Networks, Inc.

ATM cells that are received on the PE1 router are reassembled into ATM Adaptation Layer 5 (AAL5) logical link control (LLC) frames. The router removes the header and footer and adds two VLAN tags (outer and inner) to the untagged Ethernet payload based on the configuration. The VLAN IDs must correspond to the VPI and VCI of the ATM cell. You must add an MPLS label before transmitting the dual-VLAN-tagged Ethernet frame over the MPLS pseudowire. You can also add other optional MPLS tags.

NOTE: **If the AAL5 frame sent by CE1 is not encapsulated with LLC and if the untagged Ethernet payload includes the frame check sequence (FCS), PE1 rejects the AAL5 frame. Also, PE1 can transmit and receive only a dual VLAN-tagged Ethernet frame without FCS. Inclusion of FCS can result in packet drops or data corruption.**

On the PE2 router, the MPLS label and optional MPLS tags are removed and the Ethernet frame is transmitted toward the CE2 router. You can modify or remove one or both VLAN tags before forwarding the frame to the CE2 router.

Limitations

Following are the limitations of the ATM-to-Ethernet interworking feature on MX Series routers with ATM MICs:

- The ATM-to-Ethernet interworking feature is not backward compatible or does not interoperate with the ATM-Ethernet interworking feature supported on M Series and T Series Routers. The functionality is the same but the implementation is different.
- The total number of VCIs supported is 4000 for the ATM MIC. This is an existing system limit.
- If an ATM logical interface is configured with **vci-range** of *N* VCIs, then *N* VCIs are deducted from the available pool of 4000 VCIs.
- ATM quality of service (QoS) is not supported with the **vlan-vci-ccc** encapsulation. If you use the **vci-range** statement then the **vlan-vci-ccc** encapsulation supports multiple VCIs on a single logical interface. This is a hardware limitation.

Related Documentation

- Example: Configuring ATM-to-Ethernet Interworking on ATM MIC on page 295

Example: Configuring a CCC over Frame Relay Encapsulated Interface

Configure a full-duplex Layer 2 switching circuit cross-connect between Router A and Router C, using a Juniper Networks router, Router B, as the virtual switch. See the topology in Figure 17 on page 288.

Figure 17: Example Topology of a Switching Circuit Cross-Connect with Frame Relay CCC Encapsulation

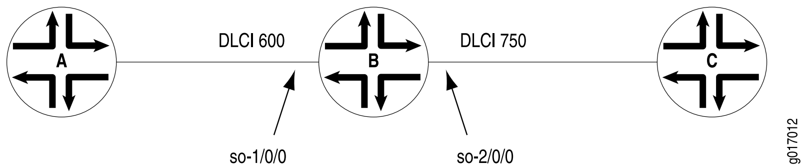


```
[edit]
interfaces {
   so-1/0/0 {
      encapsulation frame-relay-ccc;
      unit 1 {
         point-to-point;
         eui-64 frame-relay-ccc;
         dlci 600;
      }
   }
   so-2/0/0 {
      encapsulation frame-relay-ccc;
      unit 2 {
         point-to-point;
         encapsulation frame-relay-ccc;
         dlci 750;
      }
   }
}
protocols {
   connections {
      interface-switch router-a-router-c {
         interface so-1/0/0.1;
         interface so-2/0/0.2;
      }
   }
   mpls {
      interface all;
   }
}
```

Related Documentation

- *Configuring Layer 2 Switching Cross-Connects Using CCC*

Example: Configuring a TCC

Configure a full-duplex switching translational cross-connect with PPP TCC encapsulation between Router A and Router C, using a Juniper Networks router, Router B, as the virtual switch. See the topology in Figure 18 on page 289.

In this topology, Router B has a PPP connection to Router A and an ATM connection to Router C.

 Copyright © 2018, Juniper Networks, Inc.

Figure 18: Layer 2.5 Switching Translational Cross-Connect

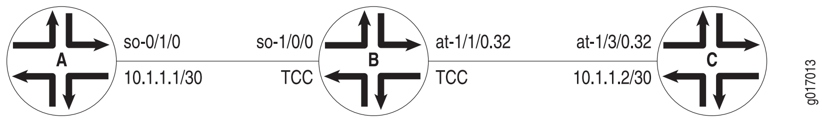

On Router A

```
[edit]
interfaces {
   so-0/1/0 {
      description "to Router B so-1/0/0";
      encapsulation ppp;
      unit 0 {
         family inet {
            address 10.1.1.1/30;
         }
      }
   }
}
```

On Router B

```
[edit]
interfaces {
   so-1/0/0 {
      description "to Router A so-0/1/0";
      encapsulation ppp-tcc;
      unit 0 {
      }
   }
   at-1/1/0 {
      description "to Router C at-0/3/0";
      atm-options {
         vpi 0 maximum-vcs 2000;
      }
      unit 32 {
         vci 32;
         encapsulation atm-tcc-vc-mux;
      }
   }
}
[edit]
protocols {
   mpls {
      interface so-1/0/0.0;
      interface at-1/1/0.32;
   }
   connections {
      interface-switch PPP-to-ATM {
         interface so-1/0/0.0;
         interface at-1/1/0.32;
      }
   }
}
```

On Router C

```
[edit]
interfaces {
```

```
    at-0/3/0 {
      description "to Router B at-1/1/0";
      atm-options {
        vpi 0 maximum-vcs 2000;
      }
      unit 32 {
        vci 32;
        encapsulation atm-vc-mux;
        family inet {
          address 10.1.1.2/30;
        }
      }
    }
  }
```

Related Documentation

- Circuit and Translational Cross-Connects Overview on page 275

Example: Configuring CCC over Aggregated Ethernet

See the topology in Figure 19 on page 290. In this topology, CE Routers A and C have aggregated Ethernet connections to PE Router B. With CCC, you specify that the circuit from Router A is connected to the circuit from Router C. Router B functions as a cross-connect switch between the two circuits. For a back-to-back connection, all VLAN IDs must be the same on Router A through Router C. You configure Router A and Router C as standard aggregated Ethernet interfaces. For more information about aggregated Ethernet, see *Aggregated Ethernet Interfaces Overview*.

Figure 19: Interface-to-Interface Circuit Cross-Connect over Aggregated Ethernet Interfaces

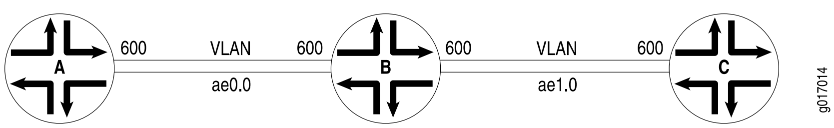

On Router A

```
[edit interfaces]
ae0 {
  vlan-tagging;
  aggregated-ether-options {
    minimum-links 1;
    link-speed 1g;
  }
  unit 0 {
    vlan-id 600;
    family inet {
      address 192.168.1.1/30;
    }
  }
}
```

On Router B

```
[edit interfaces]
ae0 {
```

 Copyright © 2018, Juniper Networks, Inc.

```
    encapsulation vlan-ccc;
    vlan-tagging;
    aggregated-ether-options {
        minimum-links 1;
        link-speed 1g;
    }
    unit 0 { # CCC switch
    encapsulation vlan-ccc;
    vlan-id 600;
    family ccc;
}
ae1 {
    encapsulation vlan-ccc;
    vlan-tagging;
    aggregated-ether-options {
        minimum-links 1;
        link-speed 100m;
    }
    unit 0 {
        encapsulation vlan-ccc;
        vlan-id 600;
        family ccc;
    }
}
[edit protocols]
mpls {
    interface all;
}
connections {
    interface-switch layer2-cross-connect {
        interface ae0.0;
        interface ae1.0;
    }
}
```

On Router C

```
[edit interfaces]
ae1 {
    vlan-tagging;
    aggregated-ether-options {
        minimum-links 1;
        link-speed 1g;
    }
    unit 0 {
        vlan-id 600;
        family inet {
            address 192.168.1.2/30;
        }
    }
}
```

Example: Configuring a Remote LSP CCC over Aggregated Ethernet

See the topology in Figure 20 on page 292. In this topology, CE Router G has an aggregated Ethernet connection to PE Router F. CE Router D has an aggregated Ethernet connection

Copyright © 2018, Juniper Networks, Inc.

to PE Router E. Router E and Router F have an MPLS LSP between them. With remote CCC, you specify that the circuit from Router D is connected to the circuit from Router G. The circuit from Router D is connected to the LSP on Router E; the circuit from Router G is connected to the LSP on Router F. In other words, **ae0.0** and **ae1.0** are connected using **lsp1-2** and **lsp2-1**. You configure Router D and Router G as standard aggregated Ethernet interfaces. For more information about aggregated Ethernet, see *Aggregated Ethernet Interfaces Overview*.

Figure 20: Remote Interface-LSP-Interface Circuit Cross-Connect over Aggregated Ethernet Interfaces

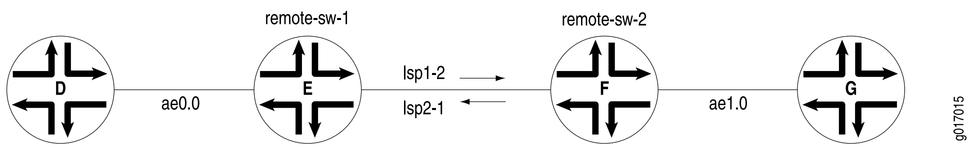

On Router D

```
[edit interface]
ae0 {
   aggregated-ether-options {
      minimum-links 1;
      link-speed 1g;
      lacp {
         active;
         periodic fast;
      }
   }
   unit 0 {
      family inet {
         address 192.168.2.1/30;
      }
   }
}
```

On Router E

```
[edit interfaces]
ae0 {
   encapsulation ethernet-ccc;
   aggregated-ether-options {
      minimum-links 1;
      link-speed 100m;
      lacp {
         active;
         periodic fast;
      }
   }
   unit 0 {
      encapsulation vlan-ccc; # default
      family ccc; # default
   }
}
[edit protocols]
mpls {
   interface all;
}
connections {
```

 Copyright © 2018, Juniper Networks, Inc.

```
    remote-interface-switch remote-sw-1 {
        interface ae0.0;
        receive-lsp lsp2_1;
        transmit-lsp lsp1_2;
    }
}
```

On Router F

```
[edit interfaces]
ae1 {
    encapsulation ethernet-ccc;
    aggregated-ether-options {
        minimum-links 1;
        link-speed 100m;
        lacp {
            active;
            periodic fast;
        }
    }
    unit 0 {
        encapsulation vlan-ccc; # default
        family ccc; # default
    }
}
[edit protocols]
mpls {
    interface all;
}
connections {
    remote-interface-switch remote-sw-2 {
        interface ae1.0;
        receive-lsp lsp1_2;
        transmit-lsp lsp2_1;
    }
}
```

On Router G

```
[edit interface]
ae1 {
    aggregated-ether-options {
        minimum-links 1;
        link-speed 1g;
        lacp {
            active;
            periodic fast;
        }
    }
    unit 0 {
        family inet {
            address 192.168.2.2/30;
        }
    }
}
```

Example: Configuring ATM-to-Ethernet Interworking

The following example shows the configuration of the ATM and Ethernet interfaces for an ATM-to-Ethernet interworking cross connect. In the example ATM DSLAM traffic is terminated on an ATM2 interface. The ATM traffic is forwarded using encapsulation type **vlan-vci-ccc** to a local Ethernet IQ2 and IQ2-E interface. See the topology in Figure 21 on page 294.

Figure 21: ATM-to-Ethernet Interworking

In this example, the ATM traffic comes from the DSLAM to the router on ATM interface **at-4/0/0** and is forwarded out on Ethernet interface **ge-2/2/1**.

```
[edit interfaces]
ge-2/2/1 {
  vlan-vci-tagging;
  encapsulation vlan-vci-ccc;
  unit 0 {
    encapsulation vlan-vci-ccc;
    vlan-id 100;
    inner-vlan-id-range start 100 end 500;
  }
}
at-4/0/0 {
  atm-options {
    vpi 100;
  }
  unit 0 {
    encapsulation vlan-vci-ccc;
    family ccc;
    vpi 100;
    vci-range start 100 end 500;
  }
}
```

Related Documentation

- Configuring ATM-to-Ethernet Interworking on page 282

 Copyright © 2018, Juniper Networks, Inc.

Example: Configuring ATM-to-Ethernet Interworking on ATM MIC

This example shows how to configure the ATM and Ethernet interfaces for an ATM-to-Ethernet interworking cross-connect.

Requirements

This example uses the following hardware and software components:

- One MX Series router with ATM MIC
- One MX Series router with Ethernet MIC
- Junos OS Release 16.1R1 or later release

Overview

Configuring ATM-to-Ethernet Interworking enables exchange of Ethernet frames between an ATM domain and an Ethernet domain on MX Series routers with ATM MIC. The ATM domain can be connected to the Ethernet domain over an MPLS pseudowire.

Topology

Consider a sample topology in which provider edge (PE) router (ATMRouter) is an MX Series router with an ATM MIC and PE router (EthernetRouter) is an MX Series router with an Ethernet MIC. CE1 and CE2 are the customer edge routers or customer-managed devices. ATMRouter and EthernetRouter are connected by means of an MPLS pseudowire. The ATM traffic between ATMRouter and CE1 comprises untagged Ethernet over ATM cells. The Ethernet traffic between EthernetRouter and CE2 comprises double-VLAN-tagged Ethernet frames.

When a packet is sent from CE1 to CE2 (ATM-to-Ethernet), ATMRouter accepts ATM cells from CE1 with virtual circuit identifier (VCI) in the range 10/50 to 10/100 and reassembles ATM cells into AAL5 frames. ATMRouter extracts the Ethernet frame from the AAL5 frame payload. ATMRouter adds two VLAN tags with VLAN IDs corresponding to the virtual path identifier (VPI) and VCI of the received ATM cell. The dual-tagged-Ethernet frame is then encapsulated into a MPLS packet and sent over the pseudowire to EthernetRouter.

EthernetRouter strips the MPLS encapsulation and the dual-VLAN-tagged Ethernet frame is sent to CE2. The outer VLAN ID is rewritten to 20 and the inner VLAN ID remains the same. The packet arrives at CE2.

The reverse happens when a packet is sent from CE2 to CE1.

Copyright © 2018, Juniper Networks, Inc.

Configuration

To enable exchange of Ethernet frames between an ATM domain and an Ethernet domain according to the topology mentioned in the overview section, perform these tasks:

- Configuring ATMRouter on page 296
- Configuring EthernetRouter on page 298

Configuring ATMRouter

Step-by-Step Procedure

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

To configure the interfaces on ATMRouter:

1. To configure the MIC to use SONET framing, include the **framing** statement at the [**edit chassis**] hierarchy level.

   ```
   [edit chassis]
   user@host# set fpc 1 pic 0 framing sonet port 0 framing sonet speed oc3-stm1
   ```

2. In configuration mode, go to the **[edit interfaces]** hierarchy level. Configure the Ethernet core interface that connects ATMRouter to EthernetRouter and specify the description of the Ethernet interface for your reference.

   ```
   [edit]
   user@host# edit interfaces
   user@host# edit ge-1/0/0
   user@host# set description PE2Facing
   ```

3. Configure a logical unit for the Ethernet interface, specify the family as **inet**, and assign an IP address to the interface. Also, specify the family as **mpls** to enable ATMRouter to connect to EthernetRouter.

   ```
   [edit interfaces ge-1/0/0]
   user@host# edit unit 0
   user@host# set family inet address 192.0.0.1/24
   user@host # set family mpls
   user@host# top
   ```

4. In configuration mode, at the **[edit interfaces]** hierarchy level, configure the ATM interface that connects to CE1 and specify the description of the ATM interface for your reference. Also, define the virtual path identifier for this interface by using the **vpi** statement and specify a value from 1 through 255. The value zero (0) is reserved and must not be used.

   ```
   [edit]
   user@host# edit interfaces
   user@host# edit at-1/2/0
   user@host# set description CE1Facing
   user@host# set atm-options vpi 10
   ```

 Copyright © 2018, Juniper Networks, Inc.

5. Configure a logical unit for the ATM interface. Also, configure the ATM logical interface by specifying the encapsulation and the protocol family. The encapsulation type **vlan-vci-vcc** supports only the **ccc** protocol family. Any attempts to configure any other interface protocol family is rejected.

```
[edit interfaces at-1/2/0]
user@host# edit unit 0
user@host# set encapsulation vlan-vci-vcc family ccc
```

6. Configure the VPI and VCI for the logical interface. The VPI value 0 is reserved and must not be used. VCI values from 0 through 31 are reserved for ATM management purposes by convention.

```
[edit interfaces at-1/2/0 unit 0]
user@host# set vpi 10 vci-range start 100 end 110
user@host# top
```

7. Configure the physical loopback interface at the [**edit interfaces**] hierarchy level.

```
[edit]
user@host# set interfaces lo0 unit 0 family inet address 198.51.100.1/32
```

8. Configure the route identifier that specifies the routing device's IP address. The router identifier is used by BGP and OSPF to identify the routing device from which a packet originated. The router identifier is usually the IP address of the local routing device. If you do not configure a router identifier, the IP address of the first interface to come online is used. This is usually the loopback interface. Otherwise, the first hardware interface with an IP address is used.

```
[edit]
user@host # set router-options router-id 198.51.100.1
```

9. At the [**edit protocols**] hierarchy level, configure the interface on which to configure MPLS as well as the loopback interface.

```
[edit protocols]
user@host# edit mpls
user@host# set interface ge-1/0/0.0
user@host# set interface lo0.0
user@host# top
```

10. At the [**edit protocols**] hierarchy level, configure a single-area OSPF network by specifying the area ID and associated interfaces (Ethernet interface and the loopback interface).

```
[edit protocols]
user@host# edit ospf
user@host# set area 0.0.0.0 interface ge-1/0/0.0
user@host# set area 0.0.0.0 interface lo0.0
user@host# top
```

Copyright © 2018, Juniper Networks, Inc.

11. Create an LDP instance on the Ethernet interface and the loopback interface at the [**edit**] hierarchy level. LDP is required as the signaling protocol for Layer 2 circuits.

    [edit]
    user@host# **edit protocols ldp**
    user@host# **set interface ge-1/0/0.0**
    user@host# **set interface lo0.0**
    user@host# **top**

12. Establish the Layer 2 circuit by specifying the **l2circuit** statement at the [**edit protocols**] hierarchy level. The neighbor parameter specifies the IP address of the PE neighbor. The interface name refers to the local CE-facing interface that forms the Layer 2 circuit. The VCI ID must match the ID of the PE neighbor.

    [edit]
    user@host# **edit protocols l2circuit**
    user@host# **edit neighbor 198.51.100.2 interface at-1/2/0.0**
    user@host# **set static incoming label 1000000 outgoing label 1000000**
    user@host# **set virtual-circuit-id 5**
    user@host# **set no-control-word**
    user@host# **set ignore-encapsulation-mismatch**
    user@host# **set ignore-mtu-mismatch**

Configuring EthernetRouter

Step-by-Step Procedure

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

To configure the interfaces on EthernetRouter:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level. Configure the Ethernet core interface that connects EthernetRouter to ATMRouter and specify the description of the Ethernet interface for your reference.

    [edit]
    user@host# **edit interfaces**
    user@host# **edit ge-1/0/0**
    user@host# **set description PE1Facing**

2. Configure a logical unit for the Ethernet interface and specify the family as **inet** and assign an IP address to the Ethernet interface. Also, specify the family as **mpls** to enable EthernetRouter to connect to ATMRouter.

    [edit interfaces ge-1/0/0]
    user@host# **edit unit 0**
    user@host# **set family inet address 192.0.0.2/24**
    user@host # **set family mpls**
    user@host# **top**

3. At the **[edit interfaces]** hierarchy level, configure the Ethernet interface that connects to CE2 and specify the description of the interface for your reference. Also, specify

 Copyright © 2018, Juniper Networks, Inc.

flexible-vlan-tagging to support transmission of 802.1Q VLAN single-tag and dual-tag frames on the same port. Specify **extended-vlan-ccc** as the encapsulation to enable tagging for translational cross-connect (TCC).

```
[edit]
user@host# edit interfaces
user@host# edit ge-1/0/1
user@host# set description CE2Facing
user@host# set flexible-vlan-tagging
user@ost# set encapsulation extended-vlan-ccc
user@host# set gigether-options ethernet-switch-profile tag-protocol-id [0x8100
   0x9100 0x88a8]
```

4. Configure a logical unit for the Ethernet interface. Also, configure mixed tagging. Mixed tagging enables you to configure two logical interfaces on the same Ethernet port, one with single-tag framing and one with dual-tag framing. You can also specify the protocol family.

```
[edit interfaces ge-1/0/1]
user@host# edit unit 0
user@host# set vlan-tags outer 0x88a8.10 inner-range 0x8100.100-110
user2host# set family ccc;
```

5. Configure the physical loopback interface at the [**edit interfaces**] hierarchy level.

```
[edit]
user@host# set interfaces lo0 unit 0 family inet address 198.51.100.2/32
```

6. Configure the route identifier that specifies the routing device’s IP address. The router identifier is used by BGP and OSPF to identify the routing device from which a packet originated. The router identifier is usually the IP address of the local routing device. If you do not configure a router identifier, the IP address of the first interface to come online is used. This is usually the loopback interface. Otherwise, the first hardware interface with an IP address is used.

```
[edit]
user@host # set router-options router-id 198.51.100.2
```

7. At the [**edit protocols**] hierarchy level, configure the interface on which to configure MPLS as well as the loopback interface.

```
[edit protocols]
user@host# edit mpls
user@host# set interface ge-1/0/0.0
user@host# set interface lo0.0
user@host# top
```

8. At the [**edit protocols**] hierarchy level, configure a single-area OSPF network by specifying the area ID and associated interfaces (Ethernet interface and the loopback interface).

```
[edit protocols]
user@host# edit ospf
```

Copyright © 2018, Juniper Networks, Inc.

```
user@host# set area 0.0.0.0 interface ge-1/0/0.0
user@host# set area 0.0.0.0 interface lo0.0
user@host# top
```

9. Create an LDP instance on the Ethernet interface and the loopback interface at the [**edit**] hierarchy level. LDP is required as the signaling protocol for Layer 2 circuits.

```
[edit]
user@host# edit protocols ldp
user@host# set interface ge-1/0/0.0
user@host# set interface lo0.0
user@host# top
```

10. Establish the Layer 2 circuit by specifying the **l2circuit** statement at the [**edit protocols**] hierarchy level. The neighbor parameter specifies the IP address of the PE neighbor. The interface name refers to the local CE-facing interface that forms the Layer 2 circuit. The VCI ID must match the ID of the PE neighbor.

```
[edit]
user@host# edit protocols l2circuit
user@host# edit neighbor 198.51.100.1 interface ge-1/0/1.0
user@host# set static incoming label 1000000 outgoing label 1000000
user@host# set virtual-circuit-id 5
user@host# set no-control-word
user@host# set ignore-encapsulation-mismatch
user@host# set ignore-mtu-mismatch
```

Related Documentation

- ATM-To-Ethernet Interworking on ATM MICs on page 286
- Verifying ATM-to-Ethernet Interworking Configuration on ATM MICs on page 300

Verifying ATM-to-Ethernet Interworking Configuration on ATM MICs

To verify that the ATM-to-Ethernet interworking feature is configured correctly, perform these tasks on both the routers:

- Verifying That The ATM Interface on Router1 Is Configured Correctly on page 301
- Verifying The Status of the MIC on Router1 on page 301
- Verify That OSPF Configuration on Router1 Is Accurate on page 301
- Verify That LDP Configuration on Router1 Is Accurate on page 302
- Verify That Layer 2 Virtual Circuit Session Configuration on Router1 Is Accurate on page 302
- Verifying That the Ethernet Interface on Router2 Is Configured Correctly on page 303
- Verifying the Status of the MIC on Router2 on page 303
- Verify That OSPF Configuration on Router2 Is Accurate on page 303

 Copyright © 2018, Juniper Networks, Inc.

- Verify That LDP Configuration on Router2 Is Accurate on page 304
- Verify That Layer 2 Virtual Circuit Session Configuration on Router2 Is Accurate on page 304

Verifying That The ATM Interface on Router1 Is Configured Correctly

Purpose To verify that the ATM interface (at-1/2/0) on Router1 is configured correctly.

Action From operational mode, enter the **show interfaces** command.

```
user@host> show interfaces at-1/2/0 terse
Interface               Admin Link Proto    Local                 Remote
 at-1/2/0          up   up
 at-1/2/0.0          up   up    ccc
 at-1/2/0.32767        up   up
```

Meaning The ATM interface on Router1 is operational.

Verifying The Status of the MIC on Router1

Purpose To verify the status of the MIC.

Action From operational mode, enter the **show chassis fpc pic-status** command.

```
user@host> show chassis fpc pic-status
Slot 0 Online
 PIC 2 Online 10x 1GE(LAN) -EH SFP
 PIC 3 Online 10x 1GE(LAN) -EH SFP
Slot 1 Online
 PIC 0 Online 2xOC12/8xOC3 CC-CE
 PIC 2 Online 10x 1GE(LAN) SFP
 PIC 3 Online 10x 1GE(LAN) SFP
Slot 2 Online
 PIC 0 Online 4x 10GE(LAN) SFP+
```

Meaning ATM MIC on FPC slot 1 is online and operational.

Verify That OSPF Configuration on Router1 Is Accurate

Purpose To verify that routers are adjacent and able to exchange OSPF data.

Action From operational mode, enter the **show ospf neighbor** command.

```
user@host> show ospf neighbor
Address          Interface              State      ID               Pri  Dead
192.0.0.1         ge-1/0/0.0              Full      198.51.100.2         128    36
```

Meaning The adjacent router is online and can accept OSPF data.

Verify That LDP Configuration on Router1 Is Accurate

Purpose To view LDP session information.

Action From operational mode, enter the **show ldp session** command.

```
user@host> show ldp session
Address State Connection Hold time Adv. Mode
 198.51.100.2    Operational Open     26 DU
```

Meaning The output indicates that the session is operational and that the connection is open. It also indicates that the session will close in 26 seconds.

Verify That Layer 2 Virtual Circuit Session Configuration on Router1 Is Accurate

Purpose To view the Layer 2 virtual circuits from the local PE router (Router1) to its neighbors.

Action From operational mode, enter the **show l2circuit connections** command.

```
user@host> show l2circuit connections
Layer-2 Circuit Connections:
Legend for connection status (St)
 EI -- encapsulation invalid NP -- interface h/w not present
 MM -- mtu mismatch Dn -- down
 EM -- encapsulation mismatch VC-Dn -- Virtual circuit Down
 CM -- control-word mismatch Up -- operational
 VM -- vlan id mismatch CF -- Call admission control failure
 OL -- no outgoing label IB -- TDM incompatible bitrate
 NC -- intf encaps not CCC/TCC TM -- TDM misconfiguration
 BK -- Backup Connection ST -- Standby Connection
 CB -- rcvd cell-bundle size bad SP -- Static Pseudowire
 LD -- local site signaled down RS -- remote site standby
 RD -- remote site signaled down HS -- Hot-standby Connection
 XX -- unknown
 Legend for interface status
 Up -- operational
 Dn -- down

Neighbor: 198.51.100.2
 Interface Type St Time last up # Up
trans
 at-1/0/0.0(vc 5) rmt Up May 24 22:01:44 2016
1
 Remote PE: 198.51.100.2, Negotiated control-word: No Encapsulation:
VLAN
 Incoming label: 299776, Outgoing label: 300192
 Negotiated PW status TLV: No
 Local interface: at-1/0/0.0, Status: Up, Encapsulation: VLAN Qin-Q
```

 Copyright © 2018, Juniper Networks, Inc.

```
and VCI Interworking
 Flow Label Transmit: No, Flow Label Receive: No
```

Meaning The command output displays the Layer 2 virtual circuits from Router1 to its neighbors.

Verifying That the Ethernet Interface on Router2 Is Configured Correctly

Purpose To verify that the Ethernet interface (ge-1/0/1) on Router2 is configured correctly.

Action From operational mode, enter the **show interfaces** command.

```
user@host> show interfaces ge-1/0/1 terse
Interface Admin Link Proto Local Remote
ge-1/0/1     up  up
ge-1/0/1.0    up  up  ccc
ge-1/0/1.32767   up  up  multiservice
```

Meaning The Ethernet interface on Router2 is operational.

Verifying the Status of the MIC on Router2

Purpose To verify the status of the MIC.

Action From operational mode, enter the **show chassis fpc pic-status** command.

```
user@host> show chassis fpc pic-status
Slot 2 Online MPC Type 1 3D Q
 PIC 0 Online 10x 1GE(LAN) SFP
 PIC 1 Online 10x 1GE(LAN) SFP
 PIC 2 Online 2xOC12/8xOC3 CC-CE
 PIC 0 Online 4x 10GE(LAN) SFP+
```

Meaning MIC 2 on MPC slot 2 is online and operational.

Verify That OSPF Configuration on Router2 Is Accurate

Purpose To verify that routers are adjacent and able to exchange OSPF data.

Action From operational mode, enter the **show ospf neighbor** command.

```
user@host> show ospf neighbor
Address          Interface             State      ID               Pri  Dead
192.0.0.1         ge-1/0/0.0             Full      198.51.100.0         128    32
```

Meaning The adjacent router is online and can accept OSPF data.

Verify That LDP Configuration on Router2 Is Accurate

Purpose To view LDP session information.

Action From operational mode, enter the **show ldp session** command.

```
user@host> show ldp session
Address State Connection Hold time Adv. Mode
 198.51.100.0    Operational Open     22  DU
```

Meaning The output indicates that the session is operational and that the connection is open. It also indicates that the session will close in 22 seconds.

Verify That Layer 2 Virtual Circuit Session Configuration on Router2 Is Accurate

Purpose To view the Layer 2 virtual circuits from the local PE router (Router2) to its neighbors.

Action From operational mode, enter the **show l2circuit connections** command.

```
user@host> show l2circuit connections
Layer-2 Circuit Connections:
Legend for connection status (St)
 EI -- encapsulation invalid NP -- interface h/w not present
 MM -- mtu mismatch Dn -- down
 EM -- encapsulation mismatch VC-Dn -- Virtual circuit Down
 CM -- control-word mismatch Up -- operational
 VM -- vlan id mismatch CF -- Call admission control failure
 OL -- no outgoing label IB -- TDM incompatible bitrate
 NC -- intf encaps not CCC/TCC TM -- TDM misconfiguration
 BK -- Backup Connection ST -- Standby Connection
 CB -- rcvd cell-bundle size bad SP -- Static Pseudowire
 LD -- local site signaled down RS -- remote site standby
 RD -- remote site signaled down HS -- Hot-standby Connection
 XX -- unknown
 Legend for interface status
 Up -- operational
 Dn -- down

Neighbor: 198.51.100.0
 Interface Type St Time last up # Up
trans
 ge-1/0/1.0(vc 5) rmt Up May 24 22:01:45 2016
1
 Remote PE: 198.51.100.0, Negotiated control-word: No Encapsulation:
VLAN Qin-Q and VCI Interworking
 Incoming label: 300192, Outgoing label: 299776                  Negotiated PW
status TLV: No
 Local interface: ge-1/0/1.0, Status: Up, Encapsulation: VLAN
 Flow Label Transmit: No, Flow Label Receive: No
```

 Copyright © 2018, Juniper Networks, Inc.

Meaning The command output displays the Layer 2 virtual circuits from Router2 to its neighbors.

Related Documentation

- ATM-To-Ethernet Interworking on ATM MICs on page 286
- Example: Configuring ATM-to-Ethernet Interworking on ATM MIC on page 295

 Copyright © 2018, Juniper Networks, Inc.

PART 2

Special Router Interfaces

Copyright © 2018, Juniper Networks, Inc.

CHAPTER 6

Configuring Discard Interfaces

Discard Interfaces Overview

The discardinterface *dsc* is not a physical interface, but a virtual interface that discards packets.

The following sections explain discard interfaces in detail:

Understanding Discard Interfaces

The discard interface allows you to identify the ingress point of a denial-of-service (DoS) attack. When your network is under attack, the target host IP address is identified, and the local policy forwards attacking packets to the discard interface. When traffic is routed out of the discard interface, the traffic is silently discarded.

The discard interface allows you to protect a network from DoS attacks by identifying the target IP address that is being attacked and configuring a policy to forward all packets to a discard interface. All packets forwarded to the discard interface are dropped. See *Example: Forwarding Packets to the Discard Interface*.

You can configure the **inet** family protocol on the discard interface, which allows you to apply an output filter to the interface. If you apply an output filter to the interface, the action specified by the filter is executed before the traffic is discarded.

Once you configure a discard interface, you must then configure a local policy to forward attacking traffic to the discard interface.

Guidelines to Follow When Configuring a Discard Interface

Keep the following guidelines in mind when configuring the discard interface:

- Only the logical interface unit 0 is supported.
- The **filter** and **address** statements are optional.

Copyright © 2018, Juniper Networks, Inc.

- Although you can configure an input filter and a filter group, these configuration statements have no effect because traffic is not transmitted from the discard interface.
- The discard interface does not support class of service (CoS).

Related Documentation

- Configuring Discard Interfaces on page 310
- *Example: Forwarding Packets to the Discard Interface*

Configuring Discard Interfaces

The discard (dsc) interface is not a physical interface, but a virtual interface that discards packets. You can configure one discard interface. This interface allows you to identify the ingress point of a denial-of-service (DoS) attack. When your network is under attack, the target host IP address is identified, and the local policy forwards attacking packets to the discard interface. Traffic routed out of the discard interface is silently discarded.

The following sections explain how to forward packets to a discard interface by configuring a discard interface with an input filter and an output filter along with an input policy to associate a community with the discard interface and an output policy to set up the community on the routes injected into the network:

Configure a Discard Interface

To configure a discard interface:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level.

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

2. Configure the discard interface.

   [edit interfaces]
   user@host# **edit dsc**

3. Configure the logical interface and the protocol family.

   [edit interfaces dsc]
   user@host# **edit unit 0 family *family***

4. Configure an input filter to be applied to received packets.

   [edit interfaces dsc unit 0 family *family*]
   user@host# **set filter input *filter-name***

5. Configure an output filter to be applied to transmitted packets.

 Copyright © 2018, Juniper Networks, Inc.

[edit interfaces dsc unit 0 family *family*]
user@host# **set filter output** ***filter-name***

6. Configure the interface IP address and the remote address of the connection.

   [edit interfaces dsc unit 0 family *family*]
   user@host# **set address** ***ip-address*** **destination** ***ip-address***

7. Commit the configuration and go to the top of the hierarchy level.

   [edit interfaces dsc unit 0 family *family*]
   user@host# **commit**
   user@host# **top**

Configure an Input Policy

You must configure an input policy to associate a community with the discard interface.

To configure an input policy to associate a community with the discard interface:

1. In configuration mode, go to the **[edit policy-options]** hierarchy level.

   [edit]
   user@host# **edit policy-options**

2. Configure the BGP community information option with a name and configure the number of community members as needed.

   [edit policy-options]
   user@host# **set community** ***community-name*** **members** ***community-id***

3. Configure a routing policy.

   [edit policy-options]
   user@host# **edit policy-statement** ***statement-name***

4. Configure a policy term with a name.

   [edit policy-options policy-statement *statement-name*]
   user@host# **edit term** ***term-variable***

5. Configure the conditions to match the source of a route with the **from** statement with the BGP community members.

   [edit policy-options policy-statement *statement-name* term *term-variable*]
   user@host# **set from community** ***value***

6. Configure the action that is to be taken when the if and to conditions match with the **then** statement. In this case, configure the next-hop address of the remote end of the point-to-point interface and accept the action.

   [edit policy-options policy-statement *statement-name* term *term-variable*]
   user@host# **set then next-hop** ***address***

Copyright © 2018, Juniper Networks, Inc.

user@host# **set then accept**

7. Commit the configuration and go to the top of the hierarchy level.

   [edit policy-options policy-statement *statement-name* term *term-variable*]
   user@host# **commit**
   user@host# **top**

Configure an Output Policy

You must configure an output policy to set up the community on the routes injected into the network.

To configure an output policy.

1. In configuration mode, go to the **[edit policy-options]** hierarchy level.

   [edit]
   user@host# **edit policy-options**

2. Configure a routing policy.

   [edit policy-options]
   user@host# **edit policy-statement** ***statement-name***

3. Configure a policy term with a name.

   [edit policy-options policy-statement *statement-name*]
   user@host# **edit term** ***term-variable***

4. Configure the list of prefix-lists of routes to match with a name.

   [edit policy-options policy-statement *statement-name* term *term-variable*]
   user@host# **set from prefix-list** ***name***

5. Configure the action that is to be taken when the if and to conditions match with the **then** statement. In this case, configure the BGP community properties (set, add, and delete) associated with a route.

   [edit policy-options policy-statement *statement-name* term *term-variable*]
   user@host# **set then community (set | add | delete)** ***community-name***

6. Commit the configuration and go to the top of the hierarchy level.

   [edit interfaces dsc unit 0 family *family*]
   user@host# **commit**
   user@host# **top**

Related Documentation

• Discard Interfaces Overview on page 309

Copyright © 2018, Juniper Networks, Inc.

CHAPTER 7

Configuring IP Demultiplexing Interfaces

Demultiplexing Interface Overview

Demultiplexing (demux) interfaces are logical interfaces that share a common, underlying interface. You can create logical subscriber interfaces using static or dynamic demultiplexing interfaces. In addition, you can use IP demultiplexing interfaces or VLAN demultiplexing interfaces when creating logical subscriber interfaces.

Demux interfaces are supported on M120 or MX Series routers only.

Demux interfaces support only Gigabit Ethernet, Fast Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet underlying interfaces.

NOTE: You can also configure demux interfaces dynamically. For information about how to configure dynamic IP demux or dynamic VLAN demux interfaces, see *Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles* or *Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles*.

To configure static demux interfaces, see "Configuring a VLAN Demultiplexing Interface" on page 320 and "Configuring an IP Demultiplexing Interface" on page 316.

IP Demux Interface Overview

IP demux interfaces use the IP source address or IP destination address to demultiplex received packets when the subscriber is not uniquely identified by a Layer 2 circuit.

Copyright © 2018, Juniper Networks, Inc.

To determine which IP demux interface to use, the destination or source prefix is matched against the destination or source address of packets that the underlying interface receives. The underlying interface family type must match the demux interface prefix type.

VLAN Demux Interface Overview

VLAN demux interfaces use the VLAN ID to demultiplex received packets when the subscriber is not uniquely identified. A VLAN demux interface uses an underlying logical interface to receive packets.

To determine which VLAN demux interface to use, the VLAN ID is matched against that which the underlying interface receives.

NOTE: **VLAN demux subscriber interfaces over aggregated Ethernet physical interfaces are supported only for MX Series routers that have only Trio MPCs installed. If the router has other MPCs in addition to Trio MPCs, theCLI accepts the configuration but errors are reported when the subscriber interfaces are brought up.**

Guidelines to Remember When Configuring A Demux Interface

Keep the following guidelines in mind when configuring the demux interface:

- Demux interfaces are supported on M120 or MX Series routers only.
- You can configure only one **demux0** interface per chassis, but you can define logical demux interfaces on top of it (for example, **demux0.1**, **demux0.2**, and so on).
- If the address in a received packet does not match any demux prefix, the packet is logically received on the underlying interface. For this reason, the underlying interface is often referred to as the *primary* interface.

- Points to Remember When Configuring an IP Demux Interface on page 314
- Points to Remember When Configuring a VLAN Demux Interface on page 315

Points to Remember When Configuring an IP Demux Interface

In addition to the guidelines in "Guidelines to Remember When Configuring A Demux Interface" on page 314, the following guidelines are to be noted when configuring an IP demux interface:

- You must associate demux interfaces with an underlying logical interface.

  NOTE: **IP demux interfaces currently support only Gigabit Ethernet, Fast Ethernet, 10-Gigabit Ethernet, and aggregated Ethernet underlying interfaces.**

- The demux underlying interface must reside on the same logical system as the demux interfaces that you configure over it.

 Copyright © 2018, Juniper Networks, Inc.

- IP demux interfaces currently supports the Internet Protocol version 4 (IPv4) suite inet and Internet Protocol version 6 (IPv6) suite inet6 family types.
- You can configure more than one demux prefix for a given demux unit. However, you cannot configure the exact same demux prefix on two different demux units with the same underlying interface.
- You can configure overlapping demux prefixes on two different demux units with the same underlying prefix. However, under this configuration, best match rules apply (in other words, the most specific prefix wins).

Points to Remember When Configuring a VLAN Demux Interface

In addition to the guidelines in "Guidelines to Remember When Configuring A Demux Interface" on page 314, the following guidelines are to be noted when configuring a VLAN demux interface:

- You must associate VLAN demux interfaces with an underlying logical interface.

NOTE: **VLAN demux interfaces currently support only Gigabit Ethernet, Fast Ethernet, 10-Gigabit Ethernet, and aggregated Ethernet underlying interfaces.**

- The demux underlying interface must reside on the same logical system as the demux interfaces that you configure over it.
- VLAN demux interfaces currently supports the Internet Protocol version 4 (IPv4) suite inet and Internet Protocol version 6 (IPv6) suite inet6 family types.

MAC Address Validation on Static Demux Interfaces

MAC address validation enables the router to validate that received packets contain a trusted IP source and an Ethernet MAC source address.

MAC address validation is supported on static demux interfaces on MX Series routers only.

There are two types of MAC address validation that you can configure:

Loose

Forwards packets when both the IP source address and the MAC source address match one of the trusted address tuples.

Drops packets when the IP source address matches one of the trusted tuples, but the MAC address does not support the MAC address of the tuple

Continues to forward packets when the source address of the incoming packet does not match any of the trusted IP addresses.

Strict

Forwards packets when both the IP source address and the MAC source address match one of the trusted address tuples.

Drops packets when the MAC address does not match the tuple's MAC source address, or when IP source address of the incoming packet does not match any of the trusted IP addresses.

Related Documentation

- *Associating VLAN IDs to VLAN Demux Interfaces*
- *Binding VLAN IDs to Logical Interfaces*
- Configuring an IP Demultiplexing Interface on page 316
- Configuring a VLAN Demultiplexing Interface on page 320
- *Subscriber Interfaces and Demultiplexing Overview*

Configuring an IP Demultiplexing Interface

Demultiplexing (demux) interfaces are logical interfaces that share a common, underlying interface. You can configure IP demultiplexing interfaces or VLAN demultiplexing interfaces.

To configure an IP demux interface, you must configure the demux prefixes that are used by the underlying interface and then configure the IP demultiplexing interface as explained in the following tasks:

1. Configuring an IP Demux Underlying Interface on page 316
2. Configuring the IP Demux Interface on page 318
3. Configuring MAC Address Validation on Static IP Demux Interfaces on page 320

Configuring an IP Demux Underlying Interface

An IP demux interface uses an underlying logical interface to receive packets. To determine which IP demux interface to use, the destination or source prefix is matched against the destination or source address of packets that the underlying interface receives. The underlying interface family type must match the demux interface prefix type.

NOTE: **IP demux interfaces currently support only Gigabit Ethernet, Fast Ethernet, 10-Gigabit Ethernet, and aggregated Ethernet underlying interfaces.**

To configure a logical interface as an IP demux underlying interface with demux source:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level:

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


Copyright © 2018, Juniper Networks, Inc.

2. Configure the interface as fe-*x/y/z* and the logical interface with the **unit** statement. Note that IP demux interfaces currently support only Gigabit Ethernet, Fast Ethernet, 10-Gigabit Ethernet, and aggregated Ethernet underlying interfaces. In this procedure, we show a Fast Ethernet interface as an example.

   [edit interfaces]
   user@host# **edit fe-*x/y/z* unit *logical-unit-number***

3. Configure the logical demux source family type on the IP demux underlying interface as inet or inet6, or both.

   [edit interfaces fe-*x/y/z* unit *logical-unit-number*]
   user@host# **set** demux-source **(inet | inet6)**

   or

   [edit interfaces fe-*x/y/z* unit *logical-unit-number*]
   user@host# **set** demux-source **[inet inet6]**

4. (Optional) To improve datapath performance for DHCPv4 subscribers, specify that only subscribers with 32-bit prefixes are allowed to come up on the interface.

   [edit interfaces fe-*x/y/z* unit *logial-unit-number*]
   user@host# **set** host-prefix-only

   NOTE: This step requires that you specify the demux-source as only **inet**. A commit error occurs if you specify only **inet6** or both **inet** and **inet6**.

5. Save the configuration and move to top of the hierarchy level.

   [edit interfaces fe-*x/y/z* unit *logial-unit-number*]
   user@host# **commit**
   user@host# **top**

To configure a logical interface as an IP demux underlying interface with demux destination:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level:

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

2. Configure the interface as fe-*x/y/z* and the logical interface with the **unit** statement. Note that IP demux interfaces currently support only Gigabit Ethernet, Fast Ethernet, 10-Gigabit Ethernet, and aggregated Ethernet underlying interfaces.

   [edit interfaces]
   user@host# **edit fe-*x/y/z* unit *logical-unit-number* unit *logical-unit-number***

3. Configure the logical demux destination family type on the IP demux underlying interface as inet or inet6.

[edit interfaces fe-*x/y/z* unit *logical-unit-number*]
user@host# **set demux-destination (inet | inet6)**

4. Save the configuration and move to top of the hierarchy level.

   [edit interfaces fe-*x/y/z* unit *logial-unit-number*]
   user@host# **commit**
   user@host# **top**

Configuring the IP Demux Interface

You can configure one or more logical demux source prefixes or destination prefixes after specifying an underlying interface for the static demux interface to use. This underlying interface must reside on the same logical system as the demux interface.

You configure demux prefixes for use by the underlying interface. The demux prefixes can represent individual hosts or networks. For a given demux interface unit, you can configure either demux source or demux destination prefixes but not both.

You can choose not to configure a demux source or demux destination prefix. This type of configuration results in a transmit-only interface.

To configure the IP demux interface with source prefix:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level:

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

2. Configure the interface as a logical demux interface (for example, demux0 interface) and configure the logical interface with the **unit** statement.

   NOTE: You can configure only one demux0 interface per chassis, but you can define logical demux interfaces on top of it (for example, demux0.1, demux0.2, and so on).

   [edit interfaces]
   user@host# **edit demux0 unit *logical-unit-number***

3. Configure the underlying interface on which the demux interface is running under the **demux-options** statement.

   [edit interfaces demux0 unit *logical-unit-number*]
   user@host# **set demux-options underlying-interface *interface-name***

4. Configure the protocol family.

   [edit interfaces demux0 unit *logical-unit-number*]
   user@host# **edit family *family***

 Copyright © 2018, Juniper Networks, Inc.

5. Configure one or more logical demux source prefixes (IP address). The prefixes are matched against the source address of packets that the underlying interface receives. When a match occurs, the packet is processed as if it was received on the demux interface.

   [edit interfaces demux0 unit *logical-unit-number* family *family*]
   user@host# **set demux-source** ***source-prefix***

6. Save the configuration and move to top of the hierarchy level.

   [edit interfaces demux0 unit *logical-unit-number* family *family*]
   user@host# **commit**
   user@host# **top**

To configure the IP demux interface with destination prefix:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level:

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

2. Configure the interface as a logical demux interface (for example, demux0 interface) and configure the logical interface with the **unit** statement.

   NOTE: You can configure only one demux0 interface per chassis, but you can define logical demux interfaces on top of it (for example, demux0.1, demux0.2, and so on).

   [edit interfaces]
   user@host# **edit demux0 unit** ***logical-unit-number***

3. Configure the underlying interface on which the demux interface is running under the **demux-options** statement.

   [edit interfaces demux0 unit *logical-unit-number*]
   user@host# **set demux-options underlying-interface** ***interface-name***

4. Configure the protocol family.

   [edit interfaces demux0 unit *logical-unit-number*]
   user@host# **edit family** ***family***

5. Configure one or more logical demux destination prefixes. The prefixes are matched against the destination address of packets that the underlying interface receives. When a match occurs, the packet is processed as if it was received on the demux interface.

   [edit interfaces demux0 unit *logical-unit-number* family *family*]
   user@host# **set demux-destination** ***destination-prefix***

6. Save the configuration and move to top of the hierarchy level.

   [edit interfaces demux0 unit *logical-unit-number* family *family*]
   user@host# **commit**
   user@host# **top**

Configuring MAC Address Validation on Static IP Demux Interfaces

MAC address validation enables the router to validate that received packets contain a trusted IP source and an Ethernet MAC source address.

To configure MAC address validation for an IP demux interface:

1. In configuration mode, go to the **[edit interfaces demux0 unit *logical-unit-number*]** hierarchy level:

   [edit]
   user@host# **edit interfaces demux0 unit *logical-unit-number***

2. Configure the protocol family for the interface.

   [edit interfaces demux0 unit *logical-unit-number*]
   user@host# **edit family *family***

3. Configure the **mac-validate** statement to validate source MAC address with loose or strict options.

   [edit interfaces demux0 unit *logical-unit-number* family *family*]
   user@host# **set mac-validate (loose | strict)**

4. Save the configuration and move to top of the hierarchy level.

   [edit interfaces demux0 unit *logical-unit-number* family *family*]
   user@host# **commit**
   user@host# **top**

Related Documentation

- Configuring a VLAN Demultiplexing Interface on page 320
- Demultiplexing Interface Overview on page 313

Configuring a VLAN Demultiplexing Interface

Demultiplexing (demux) interfaces are logical interfaces that share a common, underlying interface. You can configure IP demultiplexing interfaces or VLAN demultiplexing interfaces.

To configure a VLAN demux interface, you must configure the demux prefixes that are used by the underlying interface and then configure the VLAN demultiplexing interface as explained by the following tasks:

1. Configuring a VLAN Demux Underlying Interface on page 321
2. Configuring the VLAN Demux Interface on page 322

 Copyright © 2018, Juniper Networks, Inc.

Configuring a VLAN Demux Underlying Interface

A VLAN demux interface uses an underlying logical interface to receive packets. To determine which VLAN demux interface to use, the VLAN ID is matched against that which the underlying interface receives.

NOTE: VLAN demux interfaces currently support only Gigabit Ethernet, Fast Ethernet, 10-Gigabit Ethernet, and aggregated Ethernet underlying interfaces.

VLAN demux subscriber interfaces over aggregated Ethernet physical interfaces are supported only for MX Series routers that have only Trio MPCs installed. If the router has other MPCs in addition to Trio MPCs, the CLI accepts the configuration but errors are reported when the subscriber interfaces are brought up

To configure a logical interface as a VLAN demux underlying interface with demux source:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level:

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

2. Configure the interface as fe-*x/y/z* and the logical interface with the **unit** option.

   [edit interfaces]
   user@host# **edit fe-*x/y/z* unit *logical-unit-number* unit *logical-unit-number***

3. Configure the VLAN ID. The VLAN ID is used to determine which VLAN demux interface to use, that is the VLAN ID is matched against that which the underlying interface receives.

   [edit interfaces fe-*x/y/z* unit *logical-unit-number*]
   user@host# **set vlan-id *number***

4. Configure the logical demux source family type on the VLAN demux underlying interface as inet or inet6.

   [edit interfaces fe-*x/y/z* unit *logical-unit-number*]
   user@host# **set demux-source (inet | inet6)**

5. Save the configuration and move to top of the hierarchy level.

   [edit interfaces fe-*x/y/z* unit *logial-unit-number*]
   user@host# **commit**
   user@host# **top**

To configure a logical interface as a VLAN demux underlying interface with demux destination:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level:

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

2. Configure the interface as fe-*x/y/*z and the logical interface with the **unit** option.

   [edit interfaces]
   user@host# **edit fe-*x/y/z* unit *logical-unit-number* unit *logical-unit-number***

3. Configure the VLAN ID. The VLAN ID is used to determine which VLAN demux interface to use, that is the VLAN ID is matched against that which the underlying interface receives.

   [edit interfaces fe-*x/y/*z unit *logical-unit-number*]
   user@host# **set vlan-id *number***

4. Configure the logical demux destination family type on the VLAN demux underlying interface as inet or inet6.

   [edit interfaces fe-*x/y/*z unit *logical-unit-number*]
   user@host# **set demux-destination (inet | inet6)**

5. Save the configuration and move to top of the hierarchy level.

   [edit interfaces fe-*x/y/*z unit *logial-unit-number*]
   user@host# **commit**
   user@host# **top**

Configuring the VLAN Demux Interface

You can configure one or more logical demux source prefixes or destination prefixes after specifying an underlying interface for the static demux interface to use. This underlying interface must reside on the same logical system as the demux interface.

You configure demux prefixes for use by the underlying interface. The demux prefixes can represent individual hosts or networks. For a given demux interface unit, you can configure either demux source prefix or demux destination prefixes but not both.

You can choose not to configure a demux source prefix or a demux destination prefix. This type of configuration results in a transmit-only interface

To configure VLAN demux interface with demux source prefix:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level:

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

 Copyright © 2018, Juniper Networks, Inc.

2. Configure the interface as a logical demux interface (for example, demux0 interface) and configure the logical interface with the **unit** statement.

   NOTE: You can configure only one demux0 interface per chassis, but you can define logical demux interfaces on top of it (for example, demux0.1, demux0.2, and so on).

   [edit interfaces]
   user@host# **edit demux0 unit *logical-unit-number***

3. Configure the underlying interface on which the demux interface is running under the **demux-options** statement.

   [edit interfaces demux0 unit *logical-unit-number*]
   user@host# **set demux-options underlying-interface *interface-name***

4. Configure the protocol family for the interface.

   [edit interfaces demux0 unit *logical-unit-number*]
   user@host# **edit family *family***

5. Configure one or more logical demux source prefixes. The prefixes are matched against the source address of packets that the underlying interface receives. When a match occurs, the packet is processed as if it was received on the demux interface.

   [edit interfaces demux0 unit *logical-unit-number* family *family*]
   user@host# **set demux-source *source-prefix***

6. Save the configuration and move to top of the hierarchy level.

   [edit interfaces demux0 unit *logical-unit-number*]
   user@host# **commit**
   user@host# **top**

To configure VLAN demux interface with demux destination prefix:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level:

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

2. Configure the interface as a logical demux interface (for example, demux0 interface) and configure the logical interface with the **unit** statement.

   NOTE: You can configure only one demux0 interface per chassis, but you can define logical demux interfaces on top of it (for example, demux0.1, demux0.2, and so on).

   [edit interfaces]

user@host# **edit demux0 unit *logical-unit-number***

3. Configure the underlying interface on which the demux interface is running under the **demux-options** statement.

   [edit interfaces demux0 unit *logical-unit-number*]
   user@host# **set demux-options underlying-interface *interface-name***

4. Configure the protocol family for the interface.

   [edit interfaces demux0 unit *logical-unit-number*]
   user@host# **edit family *family***

5. Configure one or more logical demux destination prefixes. The prefixes are matched against the destination address of packets that the underlying interface receives. When a match occurs, the packet is processed as if it was received on the demux interface.

   [edit interfaces demux0 unit *logical-unit-number* family *family*]
   user@host# **set demux-destination *destination-prefix***

6. Save the configuration and move to top of the hierarchy level.

   [edit interfaces demux0 unit *logical-unit-number*]
   user@host# **commit**
   user@host# **top**

Configuring MAC Address Validation on Static VLAN Demux Interfaces

MAC address validation enables the router to validate that received packets contain a trusted IP source and an Ethernet MAC source address.

To configure MAC address validation for a VLAN demux interface:

1. In configuration mode, go to the **[edit interfaces demux0 unit *logical-unit-number*]** hierarchy level:

   [edit]
   user@host# **edit interfaces demux0 unit *logical-unit-number***

2. Configure the protocol family for the interface.

   [edit interfaces demux0 unit *logical-unit-number*]
   user@host# **edit family *family***

3. Configure the **mac-validate** statement to validate source MAC address with loose or strict options.

   [edit interfaces demux0 unit *logical-unit-number* family *family*]
   user@host# **set mac-validate (loose | strict)**

4. Save the configuration and move to top of the hierarchy level.

 Copyright © 2018, Juniper Networks, Inc.

[edit interfaces demux0 unit *logical-unit-number* family *family*]
user@host# **commit**
user@host# **top**

Verifying a Demux Interface Configuration

Purpose Check the configuration of a demux interface and its underlying interface when the following are configured:

- Two VLANs are configured, where each VLAN consists of two IP demux interfaces.
- One VLAN demultiplexes based on the source address
- The other VLAN demultiplexes based on the destination address.

Action From configuration mode on the MX Series router, run the **show interfaces fe-0/0/0** and **show interfaces demux0** configuration mode commands.

user@host> **show interfaces fe-0/0/0**

```
vlan-tagging;
unit 100 {
 vlan-id 100;
 demux-source inet; # Enable demux of inet prefixes
 family inet {
  address 10.1.1.1/24;
  filter {
   input vlan1-primary-in-filter;
   output vlan1-primary-out-filter;
  }
  mac-validate loose;
 }
}
unit 200 {
 vlan-id 200;
 demux-destination inet; # Enable demux of inet using destination addresses
 family inet {
  address 20.1.1.1/24;
 }
}
unit 300 {
 vlan-id 300;
 demux-source inet; # Enable demux of inet using source addresses
 family inet {
  address 20.1.2.1/24;
 }
}
```

user@host> **show interfaces demux0**

```
unit 101 {
 description vlan1-sub1;
 demux-options {
  underlying-interface fe-0/0/0.100;
 }
 family inet {
  demux-source 10.1.1.0/24;
```

Copyright © 2018, Juniper Networks, Inc.

```
      filter {
        input vlan1-sub1-in-filter;
        output vlan1-sub1-out-filter;
      }
    mac-validate loose;
    }
  }
  unit 102 {
    description vlan1-sub2;
    demux-options {
      underlying-interface fe-0/0/0.100;
    }
    family inet {
      demux-source {
        10.1.0.0/16;
        10.2.1.0/24;
      }
      filter {
        input vlan1-sub2-in-filter;
        output vlan1-sub2-out-filter;
      }
      mac-validate loose;
    }
  }
  unit 202 {
    description vlan2-sub2;
    demux-options {
      underlying-interface fe-0/0/0.200;
    }
    family inet {
      demux-destination 100.1.2.0/24;
    }
  }
  unit 302 {
    description vlan2-sub2;
    demux-options {
      underlying-interface fe-0/0/0.300;
    }
    family inet {
      demux-source 100.1.2.0/24;
    }
  }
```

Related Documentation

- Configuring an IP Demultiplexing Interface on page 316
- Demultiplexing Interface Overview on page 313

 Copyright © 2018, Juniper Networks, Inc.

CHAPTER 8

Configuring the Loopback Interface

Understanding the Loopback Interface

The loopback address (**lo0**) has several uses, depending on the particular Junos feature being configured. It can perform the following functions:

- Device identification—The loopback interface is used to identify the device. While any interface address can be used to determine if the device is online, the loopback address is the preferred method. Whereas interfaces might be removed or addresses changed based on network topology changes, the loopback address never changes.

  When you ping an individual interface address, the results do not always indicate the health of the device. For example, a subnet mismatch in the configuration of two endpoints on a point-to-point link makes the link appear to be inoperable. Pinging the interface to determine whether the device is online provides a misleading result. An interface might be unavailable because of a problem unrelated to the device's configuration or operation.

- Routing information—The loopback address is used by protocols such as OSPF to determine protocol-specific properties for the device or network. Further, some commands such as **ping mpls** require a loopback address to function correctly.

- Packet filtering—Stateless firewall filters can be applied to the loopback address to filter packets originating from, or destined for, the Routing Engine.

The Internet Protocol (IP) specifies a loopback network with the (IPv4) address **127.0.0.0/8**. Most IP implementations support a loopback interface (**lo0**) to represent the loopback facility. Any traffic that a computer program sends on the loopback network is addressed to the same computer. The most commonly used IP address on the loopback network is **127.0.0.1** for IPv4 and **::1** for IPv6. The standard domain name for the address is **localhost**.

The device also includes an internal loopback address (**lo0.16384**). The internal loopback address is a particular instance of the loopback address with the logical unit number 16384. Junos OS creates the loopback interface for the internal routing instance. This interface prevents any filter on **lo0.0** from disrupting internal traffic.

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- *Configuring a Loopback Interface*
- *Understanding Interfaces*
- *Understanding Management Interfaces*
- *Understanding the Discard Interface*

Configuring the Loopback Interface

Configuring the Loopback Interface

When specifying the loopback address, do not include a destination prefix. Also, in most cases, do not specify a loopback address on any unit other than unit 0.

NOTE: For Layer 3 virtual private networks (VPNs), you can configure multiple logical units for the loopback interface. This allows you to configure a logical loopback interface for each virtual routing and forwarding (VRF) routing instance. For more information, see the *Junos OS VPNs Library for Routing Devices*.

For some applications, such as SSL for Junos XML protocol, the address for the interface lo0.0 must be 127.0.0.1.

You can configure loopback interfaces using a subnetwork address for both inet and inet6 address families. Many protocols require a subnetwork address as their source address. Configuring a subnetwork loopback address as a donor interface enables these protocols to run on unnumbered interfaces.

If you configure the loopback interface, it is automatically used for unnumbered interfaces. If you do not configure the loopback interface, the router chooses the first interface to come online as the default. If you configure more than one address on the loopback interface, we recommend that you configure one to be the primary address to ensure that it is selected for use with unnumbered interfaces. By default, the primary address is used as the source address when packets originate from the interface.

For more information about unnumbered interfaces, see “Configuring an Unnumbered Interface” on page 218. For more information about primary addresses, see “Configuring the Interface Address” on page 207.

Copyright © 2018, Juniper Networks, Inc.

On the router, you can configure one physical loopback interface, **lo0**, and one or more addresses on the interface.

1. To configure the physical loopback interface, include the following statements at the **[edit interfaces]** hierarchy level:

```
[edit interfaces]
lo0 {
  unit 0 {
    family inet {
      address loopback-address;
      address <loopback-address2>;
      ...
    }
    family inet6 {
      address loopback-address;
    }
  }
}
```

Example: Configuring Two Addresses on the Loopback Interface with Host Routes

To configure two addresses on the loopback interface with host routes:

```
[edit]
user@host# edit interfaces lo0 unit 0 family inet
[edit interfaces lo0 unit 0 family inet]
user@host# set address 172.16.0.1
[edit interfaces lo0 unit 0 family inet]
user@host# set address 10.0.0.1
[edit interfaces lo0 unit 0 family inet]
user@host# top
[edit]
user@host# show
interfaces {
  lo0 {
    unit 0 {
      family inet {
        10.0.0.1;
        127.0.0.1;
        172.16.0.1;
      }
    }
  }
}
```

Example: Configuring Two Addresses on the Loopback Interface with Subnetwork Routes

To configure two addresses on the loopback interface with subnetwork routes:

```
[edit]
user@host# edit interfaces lo0 unit 0 family inet
[edit interfaces lo0 unit 0 family inet]
user@host# set address 192.16.0.1/24
[edit interfaces lo0 unit 0 family inet]
```

Copyright © 2018, Juniper Networks, Inc.

```
user@host# set address 10.2.0.1/16
[edit interfaces lo0 unit 0 family inet]
user@host# top
[edit]
user@host# show
interfaces {
  lo0 {
    unit 0 {
      family inet {
        10.2.0.1/16;
        127.0.0.1/32;
        192.16.0.1/24;
      }
    }
  }
}
```

Example: Configuring an IPv4 and an IPv6 Address on the Loopback Interface with Subnetwork Routes

To configure an IPv4 and an IPv6 address on the loopback interface with subnetwork routes:

```
[edit]
user@host# edit interfaces lo0 unit 0 family inet
[edit interfaces lo0 unit 0 family inet]
user@host# set address 192.16.0.1/24
[edit interfaces lo0 unit 0 family inet]
user@host# up
[edit interfaces lo0 unit 0 family]
user@host# edit interfaces lo0 unit 0 family inet6
[edit interfaces lo0 unit 0 family inet6]
user@host# set address 3ffe::1:200:f8ff:fe75:50df/64
[edit interfaces lo0 unit 0 family inet6]
user@host# top
[edit]
user@host# show
interfaces {
  lo0 {
    unit 0 {
      family inet {
        127.0.0.1/32;
        192.16.0.1/24;
      }
      family inet6 {
        3ffe::1:200:f8ff:fe75:50df/64;
      }
    }
  }
}
```

Related Documentation

- *Junos OS VPNs Library for Routing Devices*
- Configuring an Unnumbered Interface on page 218

 Copyright © 2018, Juniper Networks, Inc.

- Configuring the Interface Address on page 207

Copyright © 2018, Juniper Networks, Inc.

PART 3

Serial Interfaces

 Copyright © 2018, Juniper Networks, Inc.

CHAPTER 9

Serial Interfaces Overview

Serial Interfaces Overview

Devices that communicate over a serial interface are divided into two classes: data terminal equipment (DTE) and data circuit-terminating equipment (DCE). Juniper Networks Serial Physical Interface Cards (PICs) have two ports per PIC and support full-duplex data transmission. These PICs support DTE mode only. On the Serial PIC, you can configure three types of serial interfaces:

- EIA-530—An Electronics Industries Alliance (EIA) standard for the interconnection of DTE and DCE using serial binary data interchange with control information exchanged on separate control circuits.
- V.35—An ITU-T standard describing a synchronous, physical layer protocol used for communications between a network access device and a packet network. V.35 is most commonly used in the United States and in Europe.
- X.21—An ITU-T standard for serial communications over synchronous digital lines. The X.21 protocol is used primarily in Europe and Japan.

The following standards apply to serial interfaces:

- TIA/EIA Standard 530, *High-Speed 25-Position Interface for Data Terminal Equipment and Data Circuit-Terminating Equipment*, defines the signals on the cable and specifies the connector at the end of the cable.
- TIA/EIA Standard 232, *Interface between Data Terminal Equipment and Data Circuit-Terminating Equipment Employing Serial Binary Data Interchange*, describes the physical interface and protocol for serial data communication.
- ITU-T Recommendation V.35, *Data Transmission at 48 kbit/s Using 60-108 kHz Group Band Circuits*. Note that the Juniper Networks Serial PIC supports V.35 interfaces with speeds higher than 48 kilobits per second.
- ITU-T Recommendation X.21, *Interface Between Data Terminal Equipment and Data Circuit-Terminating Equipment for Synchronous Operation on Public Data Networks*.

There are no serial interface-specific logical properties. For information about general logical properties that you can configure, see *Configuring Logical Interface Properties*. This support on serial interfaces is the same as the existing LFI and MLPPP support on T1 and E1 interfaces.

Related Documentation

- Example: Physical Interface Configuration Statements for Serial Interfaces on page 337
- Configuring the Serial Line Protocol on page 338
- Configuring the Serial Clocking Mode on page 342
- Configuring the Serial Signal Handling on page 344
- Configuring the Serial DTR Circuit on page 347
- Configuring Serial Signal Polarities on page 347
- Configuring Serial Loopback Capability on page 348
- Configuring Serial Line Encoding on page 350

 Copyright © 2018, Juniper Networks, Inc.

CHAPTER 10

Configuring Serial Interfaces

Example: Physical Interface Configuration Statements for Serial Interfaces

To configure serial physical interface properties, include the **serial-options** statement at the **[edit interfaces se-** ***fpc/pic/port*****]** hierarchy level .

```
[edit interfaces se-fpc/pic/port]
serial-options {
  clock-rate rate;
  clocking-mode (dce | internal | loop);
  control-polarity (negative | positive);
  cts-polarity (negative | positive);
  dcd-polarity (negative | positive);
  dce-options {
    control-signal (assert | de-assert | normal);
    cts (ignore | normal | require);
    dcd (ignore | normal | require);
    dsr (ignore | normal | require);
    dtr signal-handling-option;
    ignore-all;
    indication (ignore | normal | require);
    rts (assert | de-assert | normal);
    tm (ignore | normal | require);
  }
  dsr-polarity (negative | positive);
  dte-options {
    control-signal (assert | de-assert | normal);
    cts (ignore | normal | require);
```

```
        dcd (ignore | normal | require);
        dsr (ignore | normal | require);
        dtr signal-handling-option;
        ignore-all;
        indication (ignore | normal | require);
        rts (assert | de-assert | normal);
        tm (ignore | normal | require);
    }
    dtr-circuit (negative | positive);
    dtr-polarity (negative | positive);
    encoding (nrz | nrzi);
    idle-cycle-flag flag;
    indication-polarity (negative | positive);
    line-protocol protocol;
    loopback mode;
    rts-polarity (negative | positive);
    tm-polarity (negative | positive);
    transmit-clock invert;
}
```

Configuring the Serial Line Protocol

Configuring the Serial Line Protocol

By default, serial interfaces use the EIA-530 line protocol. You can configure each port on the PIC independently to use one of the following line protocols:

- EIA-530
- V.35
- X.21

To configure the serial line protocol:

1. Include the **line-protocol** statement, specifying the **eia530**, **v.35**, or **x.21** option:

   ```
   line-protocol protocol;
   ```

   You can include these statements at the following hierarchy levels:

   - **[edit interfaces se-*pim*/0/*port* serial-options]**
   - **[edit interfaces se-*fpc*/*pic*/*port* serial-options]**

For more information about serial interfaces, see the following sections:

Serial Interface Default Settings

 Copyright © 2018, Juniper Networks, Inc.

Serial Interface Default Settings

EIA-530 Interface Default Settings

If you do not include the **line-protocol** statement or if you explicitly configure the default EIA-530 line protocol, the default settings are as follows:

```
dce-options | dte-options {
   cts normal;
   dcd normal;
   dsr normal;
   dtr normal;
   rts normal;
   tm normal;
}
clock-rate 16.384mhz;
clocking-mode loop;
cts-polarity positive;
dcd-polarity positive;
dsr-polarity positive;
dtr-circuit balanced;
dtr-polarity positive;
encoding nrz;
rts-polarity positive;
tm-polarity positive;
```

NOTE: On M Series routers, you can set the DCE clocking mode for EIA-530 interfaces and commit. An error message is not displayed and the CLI is not blocked.

You can include the line-protocol statement at the following hierarchy levels:

- **[edit interfaces se-*pim*/0/*port* serial-options]**
- **[edit interfaces se-*fpc*/*pic*/*port* serial-options]**

V.35 Interface Default Settings

If you include the **line-protocol v.35** statement, the default settings are as follows:

```
dce-options | dte-options {
   cts normal;
   dcd normal;
   dsr normal;
   dtr normal;
   rts normal;
}
clock-rate 16.384mhz;
clocking-mode loop;
```

Copyright © 2018, Juniper Networks, Inc.

```
cts-polarity positive;
dcd-polarity positive;
dsr-polarity positive;
dtr-circuit balanced;
dtr-polarity positive;
encoding nrz;
rts-polarity positive;
```

You can include the line-protocol statement at the following hierarchy levels:

- **[edit interfaces se-*pim*/0/*port* serial-options]**
- **[edit interfaces se-*fpc*/*pic*/*port* serial-options]**

X.21 Interface Default Settings

If you include the **line-protocol x.21** statement, the default settings are as follows:

```
dce-options | dte-options {
  control-signal normal;
  indication normal;
}
clock-rate 16.384mhz;
clocking-mode loop;
control-polarity positive;
encoding nrz;
indication-polarity positive;
```

You can include the line-protocol statement at the following hierarchy levels:

- **[edit interfaces se-*pim*/0/*port* serial-options]**
- **[edit interfaces se-*fpc*/*pic*/*port* serial-options]**

Invalid Serial Interface Statements

The following sections show the invalid configuration statements for each type of serial interface. If you include the following statements in the configuration, an error message indicates the location of the error and the configuration is not activated.

- Invalid EIA-530 Interface Statements on page 340
- Invalid V.35 interface Statements on page 341
- Invalid X.21 Interface Statements on page 341

Invalid EIA-530 Interface Statements

If you do not include the **line-protocol** statement or if you explicitly configure the default EIA-530 line protocol, the following statements are invalid:

```
dce-options | dte-options {
  control-signal (assert | de-assert | normal);
  indication (ignore | normal | require);
}
control-polarity (negative | positive);
indication-polarity (negative | positive);
```

You can include the line-protocol statement at the following hierarchy levels:

 Copyright © 2018, Juniper Networks, Inc.

- **[edit interfaces se-*pim*/0/*port* serial-options]**
- **[edit interfaces se-*fpc*/*pic*/*port* serial-options]**

Invalid V.35 interface Statements

If you include the **line-protocol v.35** statement, the following statements are invalid:

```
dce-options | dte-options {
  control-signal (assert | de-assert | normal);
  indication (ignore | normal | require);
  tm (ignore | normal | require);
}
control-polarity (negative | positive);
indication-polarity (negative | positive);
loopback (dce-local | dce-remote);
tm-polarity (negative | positive);
```

You can include the line-protocol statement at the following hierarchy levels:

- **[edit interfaces se-*pim*/0/*port* serial-options]**
- **[edit interfaces se-*fpc*/*pic*/*port* serial-options]**

Invalid X.21 Interface Statements

If you include the **line-protocol x.21** statement, the following statements are invalid:

```
dce-options | dte-options {
  cts (ignore | normal | require);
  dcd (ignore | normal | require);
  dsr (ignore | normal | require);
  dtr (assert | de-assert | normal);
  rts (assert | de-assert | normal);
  tm (ignore | normal | require);
}
clocking-mode (dce | internal);
cts-polarity (negative | positive);
dce-polarity (negative | positive);
dsr-polarity (negative | positive);
dtr-circuit (balanced | unbalanced);
dtr-polarity (negative | positive);
loopback (dce-local | dce-remote);
rts-polarity (negative | positive);
tm-polarity (negative | positive);
```

You can include the line-protocol statement at the following hierarchy levels:

- **[edit interfaces se-*pim*/0/*port* serial-options]**
- **[edit interfaces se-*fpc*/*pic*/*port* serial-options]**

See Also Serial Interfaces Overview on page 335

Configuring the Serial Clocking Mode

Configuring the Serial Clocking Mode

By default, serial interfaces use loop clocking mode. For EIA-530 and V.35 interfaces, you can configure each port on the PIC independently to use loop, DCE, or internal clocking mode. For X.21 interfaces, only loop clocking mode is supported.

The three clocking modes work as follows:

- Loop clocking mode—Uses the DCE's RX clock to clock data from the DCE to the DTE.
- DCE clocking mode—Uses the TXC clock, which is generated by the DCE specifically to be used by the DTE as the DTE's transmit clock.
- Internal clocking mode—Also known as line timing, uses an internally generated clock. You can configure the speed of this clock by including the **clock-rate** statement at the **[edit interfaces se-*pim*/0/*port* serial-options]** or **[edit interfaces se-*fpc*/*pic*/*port* dte-options]** hierarchy levels. For more information about the DTE clock rate, see "Configuring the DTE Clock Rate" on page 343.

Note that DCE clocking mode and loop clocking mode use external clocks generated by the DCE.

Figure 22 on page 342 shows the clock sources of loop, DCE, and internal clocking modes.

Figure 22: Serial Interface Clocking Mode

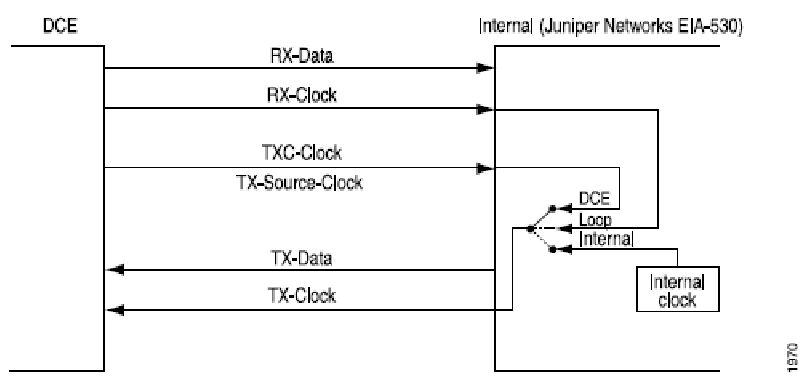

To configure the clocking mode of a serial interface, include the **clocking-mode** statement:

```
clocking-mode (dce | internal | loop);
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces se-*pim*/0/*port* serial-options]**

 Copyright © 2018, Juniper Networks, Inc.

- **[edit interfaces se-*fpc/pic/port* serial-options]**

Inverting the Serial Interface Transmit Clock

When an externally timed clocking mode (DCE or loop) is used, long cables might introduce a phase shift of the DTE-transmitted clock and data. At high speeds, this phase shift might cause errors. Inverting the transmit clock corrects the phase shift, thereby reducing error rates.

By default, the transmit clock is not inverted. To invert the transmit clock, include the **transmit-clock invert** statement:

```
transmit-clock invert;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces se-*pim/0/port* serial-options]**
- **[edit interfaces se-*fpc/pic/port* serial-options]**

Configuring the DTE Clock Rate

By default, the serial interface has a clock rate of 16.384 MHz. For EIA-530 and V.35 interfaces with internal clocking mode configured, you can configure the clock rate.

To configure the clock rate, include the **clock-rate** statement:

```
clock-rate rate;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces se-*pim/0/port* serial-options]**
- **[edit interfaces se-*fpc/pic/port* serial-options]**

You can configure the following interface speeds:

- 2.048 MHz
- 2.341 MHz
- 2.731 MHz
- 3.277 MHz
- 4.096 MHz
- 5.461 MHz
- 8.192 MHz
- 16.384 MHz

Although the serial interface is intended for use at the default rate of 16.384 MHz, you might need to use a slower rate if any of the following conditions prevail:

Copyright © 2018, Juniper Networks, Inc.

- The interconnecting cable is too long for effective operation.
- The interconnecting cable is exposed to an extraneous noise source that might cause an unwanted voltage in excess of +1 volt measured differentially between the signal conductor and circuit common at the load end of the cable, with a 50-ohm resistor substituted for the generator.
- You need to minimize interference with other signals.
- You need to invert signals.

For detailed information about the relationship between signaling rate and interface cable distance, see the following standards:

- EIA-422-A, *Electrical Characteristics of Balanced Voltage Digital Interface Circuits*
- EIA-423-A, *Electrical Characteristics of Unbalanced Voltage Digital Interface Circuits*

Related Documentation

- Serial Interfaces Overview on page 335

Configuring the Serial Signal Handling

By default, normal signal handling is enabled for all signals. For each signal, the **normal** option applies to the normal signal handling for that signal, as defined by the following standards:

- TIA/EIA Standard 530
- ITU-T Recommendation V.35
- ITU-T Recommendation X.21

Table 50 on page 344 shows the serial interface modes that support each signal type.

Table 50: Signal Handling by Serial Interface Type

Signal	Serial Interfaces
From-DCE signals	
Clear to send (CTS)	EIA-530 and V.35
Data carrier detect (DCD)	EIA-530 and V.35
Data set ready (DSR)	EIA-530 and V.35
Indication	X.21 only
Test mode (TM)	EIA-530 only
To-DCE signals	
Control signal	X.21 only

Copyright © 2018, Juniper Networks, Inc.

Table 50: Signal Handling by Serial Interface Type (continued)

Signal	Serial Interfaces
Data transfer ready (DTR)	EIA-530 and V.35
Request to send (RTS)	EIA-530 and V.35

You configure serial interface signal characteristics by including the **dce-options** or **dte-options** statement:

```
dce-options |dte-options {
  control-signal (assert | de-assert | normal);
  cts (ignore | normal | require);
  dcd (ignore | normal | require);
  dsr (ignore | normal | require);
  dtr signal-handling-option;
  ignore-all;
  indication (ignore | normal | require);
  rts (assert | de-assert | normal);
  tm (ignore | normal | require);
}
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces se-*pim*/0/*port* serial-options]**
- **[edit interfaces se-*fpc*/*pic*/*port* serial-options]**

For EIA-530 and V.35 interfaces, configure to-DCE signals by including the **dtr** and **rts** statements, specifying the **assert**, **de-assert**, or **normal** option:

```
dtr (assert | de-assert | normal);
rts (assert | de-assert | normal);
```

For X.21 interfaces, configure to-DCE signals by including the **control-signal** statement, specifying the **assert**, **de-assert**, or **normal** option:

```
control-signal (assert | de-assert | normal);
```

Assertion is when the positive side of a given signal is at potential high-level output voltage (Voh), while the negative side of the same signal is at potential low-level output voltage (Vol). *Deassertion* is when the positive side of a given signal is at potential Vol, while the negative side of the same signal is at potential Voh.

For the DTR signal, you can configure normal signal handling using the signal for automatic resynchronization by including the **dtr** statement, and specifying the **auto-synchronize** option:

```
dtr {
  auto-synchronize {
    duration milliseconds;
    interval seconds;
  }
}
```

Copyright © 2018, Juniper Networks, Inc.

The pulse duration of resynchronization can be from 1 through 1000 milliseconds. The offset interval for resynchronization can be from 1 through 31 seconds.

For EIA-530 and V.35 interfaces, configure from-DCE signals by including the **cts**, **dcd**, and **dsr** statements, specifying the **ignore**, **normal**, or **require** option:

```
cts (ignore | normal | require);
dcd (ignore | normal | require);
dsr (ignore | normal | require);
```

For X.21 interfaces, configure from-DCE signals by including the **indication** statement, specifying the **ignore**, **normal**, or **require** option:

```
indication (ignore | normal | require);
```

For EIA-530 interfaces only, you can configure from-DCE test-mode (TM) signaling by including the **tm** statement, specifying the **ignore**, **normal**, or **require** option:

```
tm (ignore | normal | require);
```

To specify that the from-DCE signal must be asserted, include the **require** option in the configuration. To specify that the from-DCE signal must be ignored, include the **ignore** option in the configuration.

NOTE: For V.35 and X.21 interfaces, you cannot include the **tm** statement in the configuration.

For X.21 interfaces, you cannot include the **cts**, **dcd**, **dsr**, **dtr**, and **rts** statements in the configuration.

For EIA-530 and V.35 interfaces, you cannot include the **control-signal** and **indication** statements in the configuration.

For a complete list of serial options statements that are not supported by each serial interface mode, see "Invalid Serial Interface Statements" on page 340.

To return to the default normal signal handling, delete the **require**, **ignore**, **assert**, **de-assert**, or **auto-synchronize** statement from the configuration, as shown in the following example:

```
[edit]
user@host# delete interfaces se-fpc/pic/port dte-options control-leads cts require
```

To explicitly configure normal signal handling, include the **control-signal** statement with the **normal** option:

```
control-signal normal;
```

You can configure the serial interface to ignore all control leads by including the **ignore-all** statement:

```
ignore-all;
```


Copyright © 2018, Juniper Networks, Inc.

You can include the **ignore-all** statement in the configuration only if you do not explicitly enable other signal handling options at the **[edit interfaces se-*pim*/0/*port* serial-options dce-options]** or **[edit interfaces se-*fpc*/*pic*/*port* serial-options dte-options]** hierarchy levels.

You can include the **control-signal**, **cts**, **dcd**, **dsr**, **dtr**, **indication**, **rts**, and **tm** statements at the following hierarchy levels:

- **[edit interfaces se-*pim*/0/*port* serial-options dte-options]**
- **[edit interfaces se-*fpc*/*pic*/*port* serial-options dte-options]**

Configuring the Serial DTR Circuit

A balanced circuit has two currents that are equal in magnitude and opposite in phase. An unbalanced circuit has one current and a ground; if a pair of terminals is unbalanced, one side is connected to electrical ground and the other carries the signal. By default, the DTR circuit is balanced.

For EIA-530 and V.35 interfaces, configure the DTR circuit by including the **dtr-circuit** statement:

```
dtr-circuit (balanced | unbalanced);
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces se-*pim*/0/*port* serial-options]**
- **[edit interfaces se-*fpc*/*pic*/*port* serial-options]**

Configuring Serial Signal Polarities

Serial interfaces use a differential protocol signaling technique. Of the two serial signals associated with a circuit, the one referred to as the A signal is denoted with a plus sign, and the one referred to as the B signal is denoted with a minus sign; for example, DTR+ and DTR–. If DTR is low, then DTR+ is negative with respect to DTR–. If DTR is high, then DTR+ is positive with respect to DTR–.

By default, all signal polarities are positive. You can reverse this polarity on a Juniper Networks serial interface. You might need to do this if signals are miswired as a result of reversed polarities.

For EIA-530 and V.35 interfaces, configure signal polarities by including the **cts-polarity**, **dcd-polarity**, **dsr-polarity**, **dtr-polarity**, **rts-polarity**, and **tm-polarity** statements:

```
cts-polarity (negative | positive);
dcd-polarity (negative | positive);
dsr-polarity (negative | positive);
dtr-polarity (negative | positive);
rts-polarity (negative | positive);
tm-polarity (negative | positive);
```

You can include these statements at the following hierarchy levels:

Copyright © 2018, Juniper Networks, Inc.

- [edit interfaces se-*pim*/0/*port* serial-options]
- [edit interfaces se-*fpc*/*pic*/*port* serial-options]

For X.21 interfaces, configure signal polarities by including the **control-polarity** and **indication-polarity** statements:

```
control-polarity (negative | positive);
indication-polarity (negative | positive);
```

You can include these statements at the following hierarchy levels:

- [edit interfaces se-*pim*/0/*port* serial-options]
- [edit interfaces se-*fpc*/*pic*/*port* serial-options]

Configuring Serial Loopback Capability

From the router, remote line interface unit (LIU) loopback loops the TX (transmit) data and TX clock back to the router as RX (receive) data and RX clock. From the line, LIU loopback loops the RX data and RX clock back out the line as TX data and TX clock, as shown in Figure 23 on page 348.

Figure 23: Serial Interface LIU Loopback

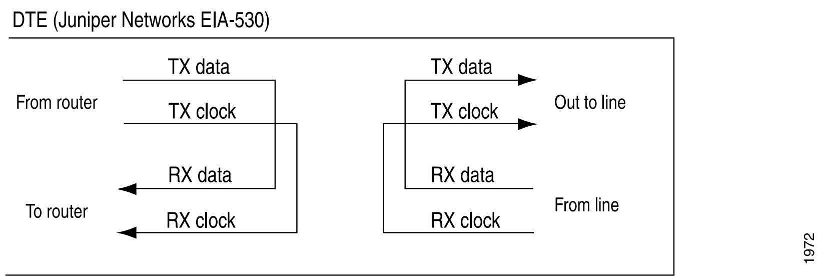

DCE local and DCE remote control the EIA-530 interface-specific signals for enabling local and remote loopback on the link partner DCE. Local loopback is shown in Figure 24 on page 348.

Figure 24: Serial Interface Local Loopback

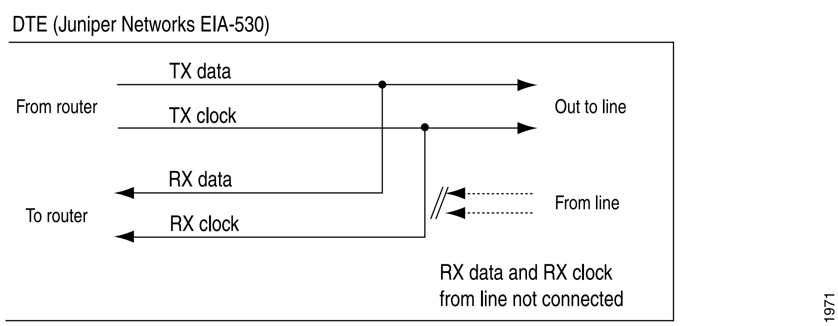

For EIA-530 interfaces, you can configure DCE local, DCE remote, local, and remote (LIU) loopback capability.

Copyright © 2018, Juniper Networks, Inc.

For V.35, you can configure remote LIU and local loopback capability. DCE local and DCE remote loopbacks are not supported on V.35 and X.21 interfaces. Local and remote loopbacks are not supported on X.21 interfaces.

To configure the loopback capability on a serial interface, include the **loopback** statement, specifying the **dce-local**, **dce-remote**, **local**, or **remote** option:

```
loopback mode;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces se-*pim*/0/*port* serial-options]**
- **[edit interfaces se-*fpc*/*pic*/*port* serial-options]**

To disable the loopback capability, remove the **loopback** statement from the configuration:

```
[edit]
user@host# delete interfaces se-fpc/pic/port serial-options loopback
```

You can determine whether there is an internal or external problem by checking the error counters in the output of the **show interface se-*fpc*/*pic*/*port* extensive** command:

```
user@host> show interfaces se-fpc/pic/port extensive
```

To Configure Serial Loopback Capability:

1. To determine the source of a problem, loop the packets on the local router, the local DCE, the remote DCE, and the remote line interface unit (LIU).

2. To do this, include the **no-keepalives** and **encapsulation cisco-hdlc** statements at the **[edit interfaces se-*fpc*/*pic*/*port*]** hierarchy level, and the **loopback local** option at the **[edit interfaces se-*pim*/0/*port* serial-options]** or **[edit interfaces se-*fpc*/*pic*/*port* serial-options]** hierarchy level. With this configuration, the link stays up, so you can loop ping packets to a remote router. The **loopback local** statement causes the interface to loop within the PIC just before the data reaches the transceiver.

   ```
   [edit interfaces]
   se-1/0/0 {
      no-keepalives;
      encapsulation cisco-hdlc;
      serial-options {
         loopback local;
      }
      unit 0 {
         family inet {
            address 10.100.100.1/24;
         }
      }
   }
   ```

Related Documentation

- Serial Interfaces Overview on page 335

Configuring Serial Line Encoding

By default, serial interfaces use non-return to zero (NRZ) line encoding. You can configure non-return to zero inverted (NRZI) line encoding if necessary.

To have the interface use NRZI line encoding, include the **encoding** statement, specifying the **nrzi** option:

```
encoding nrzi;
```

To explicitly configure the default NRZ line encoding, include the **encoding** statement, specifying the **nrz** option:

```
encoding nrz;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces se-*pim*/0/*port* serial-options]**
- **[edit interfaces se-*fpc*/*pic*/*port* serial-options]**

When setting the line encoding parameter, you must set the same value for paired ports. Ports 0 and 1 must share the same value.

Specifying a USB Modem Interface on J Series Routers

The J Series routers contain two USB ports controlled by a single USB controller. One USB port can support USB devices, while the other one can act as a USB modem.

The USB modem provides a dial-in remote management interface, and supports dialer interface features by sharing the same dial pool as a dialer interface. The dial pool allows the logical dialer interface (**dl*n***) and the physical interface (**umd0**) to be bound together dynamically on a per-call basis.

The following dialer interface features are supported by the USB modem interface:

- Encapsulation PPP
- CoS
- NAT
- Interface statistics
- Packet capture
- GRE tunnel
- Stateful firewall
- Traffic sampling

To configure a USB modem interface, include the following statements at the **[edit interfaces]** hierarchy level:

```
[edit interfaces]
```

 Copyright © 2018, Juniper Networks, Inc.

```
umd0 {
   dialer-options {
      pool pool-name <priority priority>;
   }
   modem-options {
      dialin (console | routable);
      init-command-string initialization-command-string;
   }
}
```

The pool name specified at the **[edit interfaces umd0 dialer-options pool]** hierarchy level must be the same as the pool name specified at the **[edit interfaces dl*n* unit *logical-unit-number* dialer-options pool]** hierarchy level.

Configure the USB modem to operate as a dial-in WAN backup interface by including the **dialin** statement and specifying the **routable** option. If the USB modem is to be used as a dial-in console, specify the **console** option in the **dialin** statement.

When the Services Router applies the modem AT commands configured in the **init-command-string** statement or the default sequence of initialization commands to the modem, it compares them to the initialization commands already configured on the modem and makes the following changes:

- If the commands are the same, the router overrides the existing modem values that do not match. For example, if the initialization commands on the modem include S0=0 and the router's **init-command-string** configuration includes S0=2, the Services Router applies S0=2.
- If the initialization commands on the modem do not include a command in the router's **init-command-string** statement configuration, the router adds it. For example, if the **init-command-string** statement includes the command L2, but the modem commands do not include it, the router adds L2 to the initialization commands configured on the modem.

Include the following statements at the **[edit interfaces dl*n*]** hierarchy level to support a minimum configuration for a dialer interface connected to a USB modem:

```
[edit interfaces dln]
encapsulation ppp;
unit logical-unit-number;
dialer-options {
   dial-string dial-string-numbers;
   pool pool-name <priority priority>;
}
ppp-options {
   chap;
   access-profile name;
   local-name name;
   passive;
}
family inet {
   mtu bytes;
   address address {
      destination address;
```

```
  }
}
```

 Copyright © 2018, Juniper Networks, Inc.

PART 4

Monitoring and Troubleshooting Interfaces

 Copyright © 2018, Juniper Networks, Inc.

CHAPTER 11

Monitoring Interfaces

Tracing Interface Operations Overview

You can trace the operations of individual router interfaces and those of the interface process (dcd). For a general discussion of tracing and of the precedence of multiple tracing operations, see the *Junos OS Administration Library*.

For information about the operations of Virtual Router Resolution Protocol (VRRP)-enabled interfaces, see the *Junos OS High Availability Library for Routing Devices*.

Related Documentation

- Tracing Operations of an Individual Router Interface on page 355
- Tracing Operations of the Interface Process on page 356

Tracing Operations of an Individual Router Interface

To trace the operations of individual router interfaces, perform the following steps:

1. In configuration mode, go to the **[edit interfaces *interface-name*]** hierarchy level:

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

2. Configure the **traceoptions** option.

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

3. Configure the tracing flag.

   [edit interfaces *interface-name* traceoptions]
   user@host# **set flag *flag-option***

You can specify the following interface tracing flags:

- **all**—Trace all interface operations.
- **event**—Trace all interface events.
- **ipc**—Trace all interface interprocess communication (IPC) messages.
- **media**—Trace all interface media changes.

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** files.

For more information about trace operations, see "Tracing Operations of the Interface Process" on page 356.

Related Documentation

- traceoptions on page 1019

Tracing Operations of the Interface Process

To trace the operations of the router or switch interface process, dcd, perform the following steps:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level:

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

2. Configure the **traceoptions** statement.

   [edit interfaces]
   user@host# **edit traceoptions**

3. Configure the **no-remote-trace** option to disable remote tracing.

   [edit interfaces traceoptions]
   user@host# **set no-remote-trace**

4. Configure the **file** ***filename*** option.

   [edit interfaces traceoptions]
   user@host# **edit file**

5. Configure the **files** ***number*** option, **match** ***regular-expression*** option, **size** ***size*** option, and **world-readable | no-world-readable** option.

   [edit interfaces traceoptions file]
   user@host# **set files** ***number***
   user@host# **set match** ***regular-expression***
   user@host# **set size** ***size***
   user@host# **set word-readable | no-world-readable**

6. Configure the tracing flag.

Copyright © 2018, Juniper Networks, Inc.

[edit interfaces traceoptions]
user@host# **set flag *flag-option***

7. Configure the **disable** option in **flag *flag-option*** statement to disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

   [edit interfaces traceoptions]
   user@host# **set flag *flag-option* disable**

You can specify the following flags in the **interfaces traceoptions** statement:

- **all**—Enable all configuration logging.
- **change-events**—Log changes that produce configuration events.
- **gres-events**—Log the events related to GRES.
- **resource-usage**—Log the resource usage for different states.
- **config-states**—Log the configuration state machine changes.
- **kernel**—Log configuration IPC messages to kernel.
- **kernel-detail**—Log details of configuration messages to kernel.
- **select-events**—Log the events on select state machine.

By default, interface process operations are placed in the file named dcd and three 1-MB files of tracing information are maintained.

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

Related Documentation

- Tracing Interface Operations Overview on page 355
- Tracing Operations of an Individual Router Interface on page 355
- **traceoptions on page 1019**

Monitoring a PPP Session

You can monitor PPP packet exchanges. When monitoring is enabled, packets exchanged during a session are logged by default to **/var/log/pppd**, or to the file specified in the **traceoptions** statement.

To monitor a PPP session:

1. In configuration mode, go to the **[edit protocols ppp]** hierarchy level.

   [edit]
   user@host# **edit protocols ppp**

2. Include the **monitor-session** statement.

[edit protocols ppp]
user@host# **monitor-session (*interface-name* | all);**

When monitoring is configured, the operational mode commands **show ppp summary** and **show ppp interface** display a **Monitored** flag in the **Session flags** column or line.

Related Documentation monitor-session on page 782

Tracing Operations of the pppd Process

You can trace the operations of the router's pppd process.

To trace the router's pppd process:

1. In configuration mode, go to the **[edit protocols ppp]** hierarchy level.

   [edit]
   user@host# **edit protocols ppp**

2. Include the **traceoptions** statement.

   [edit protocols ppp]
   traceoptions {
   file *filename* <files *number*> <match *regular-expression*> <size *size*> <world-readable | no-world-readable>;
   flag *flag*;
   level *severity-level*;
   no-remote-trace;
   }

- To specify more than one tracing operation, include multiple **flag** statements.

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

  - **access**—Trace access code
  - **address-pool**—Trace address pool code
  - **all**—Trace all areas of code
  - **auth**—Trace authentication code
  - **chap**—Trace challenge handshake authentication protocol code
  - **ci**—Trace CI code
  - **config**—Trace configuration code
  - **ifdb**—Trace interface database code
  - **lcp**—Trace LCP state machine code
  - **memory**—Trace memory management code
  - **message**—Trace message processing code

 Copyright © 2018, Juniper Networks, Inc.

- **mlppp**—Trace multilink point-to-point protocol code
- **ncp**—Trace NCP state machine code
- **pap**—Trace password authentication protocol code
- **ppp**—Trace PPP protocol processing code
- **radius**—Trace RADIUS processing code
- **redundancy**—Trace redundancy code
- **rtsock**—Trace routing socket code
- **session**—Trace session management code
- **signal**—Trace signal handling code
- **timer**—Trace timer code
- **ui**—Trace user interface code

Related Documentation

traceoptions on page 1033

Copyright © 2018, Juniper Networks, Inc.

CHAPTER 12

Troubleshooting Interfaces

Configuring Interface Diagnostics Tools to Test the Physical Layer Connections

Configuring Loopback Testing

Loopback testing allows you to verify the connectivity of a circuit. You can configure any of the following interfaces to execute a loopback test: aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, E1, E3, *N*xDS0, serial, SONET/SDH, T1, and T3.

The physical path of a network data circuit usually consists of segments interconnected by devices that repeat and regenerate the transmission signal. The transmit path on one device connects to the receive path on the next device. If a circuit fault occurs in the form of a line break or a signal corruption, you can isolate the problem by using a loopback test. Loopback tests allow you to isolate segments of the circuit and test them separately.

To do this, configure a *line loopback* on one of the routers. Instead of transmitting the signal toward the far-end device, the line loopback sends the signal back to the originating router. If the originating router receives back its own Data Link Layer packets, you have verified that the problem is beyond the originating router. Next, configure a line loopback farther away from the local router. If this originating router does not receive its own Data Link Layer packets, you can assume that the problem is on one of the segments between the local router and the remote router's interface card. In this case, the next troubleshooting step is to configure a line loopback closer to the local router to find the source of the problem.

Copyright © 2018, Juniper Networks, Inc.

The following types of loopback testing are supported by Junos OS:

- DCE local—Loops packets back on the local data circuit-terminating equipment (DCE).
- DCE remote—Loops packets back on the remote DCE.
- Local—Useful for troubleshooting physical PIC errors. Configuring local loopback on an interface allows transmission of packets to the channel service unit (CSU) and then to the circuit toward the far-end device. The interface receives its own transmission, which includes data and timing information, on the local router's PIC. The data received from the CSU is ignored. To test a local loopback, issue the **show interfaces *interface-name*** command. If PPP keepalives transmitted on the interface are received by the PIC, the **Device Flags** field contains the output **Loop-Detected**.
- Payload—Useful for troubleshooting the physical circuit problems between the local router and the remote router. A payload loopback loops data only (without clocking information) on the remote router's PIC. With payload loopback, overhead is recalculated.
- Remote—Useful for troubleshooting the physical circuit problems between the local router and the remote router. A remote loopback loops packets, including both data and timing information, back on the remote router's interface card. A router at one end of the circuit initiates a remote loopback toward its remote partner. When you configure a remote loopback, the packets received from the physical circuit and CSU are received by the interface. Those packets are then retransmitted by the PIC back toward the CSU and the circuit. This loopback tests all the intermediate transmission segments.

Table 51 on page 362 shows the loopback modes supported on the various interface types.

Table 51: Loopback Modes by Interface Type

Interface	Loopback Modes	Usage Guidelines
Aggregated Ethernet, Fast Ethernet, Gigabit Ethernet	Local	*Configuring Ethernet Loopback Capability*
Circuit Emulation E1	Local and remote	*Configuring E1 Loopback Capability*
Circuit Emulation T1	Local and remote	*Configuring T1 Loopback Capability*
E1 and E3	Local and remote	*Configuring E1 Loopback Capability* and *Configuring E3 Loopback Capability*
*N*xDS0	Payload	*Configuring NxDS0 IQ and IQE Interfaces*, *Configuring T1 and NxDS0 Interfaces*, *Configuring Channelized OC12/STM4 IQ and IQE Interfaces (SONET Mode)*, *Configuring Fractional E1 IQ and IQE Interfaces*, and *Configuring Channelized T3 IQ Interfaces*
Serial (V.35 and X.21)	Local and remote	"Configuring Serial Loopback Capability" on page 348

Copyright © 2018, Juniper Networks, Inc.

Table 51: Loopback Modes by Interface Type (continued)

Interface	Loopback Modes	Usage Guidelines
Serial (EIA-530)	DCE local, DCE remote, local, and remote	"Configuring Serial Loopback Capability" on page 348
SONET/SDH	Local and remote	*Configuring SONET/SDH Loopback Capability to Identify a Problem as Internal or External*
T1 and T3	Local, payload, and remote	*Configuring T1 Loopback Capability* and *Configuring T3 Loopback Capability* See also *Configuring the T1 Remote Loopback Response*

To configure loopback testing, include the **loopback** statement:

user@host# loopback *mode*;

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* aggregated-ether-options]**
- **[edit interfaces *interface-name* ds0-options]**
- **[edit interfaces *interface-name* e1-options]**
- **[edit interfaces *interface-name* e3-options]**
- **[edit interfaces *interface-name* fastether-options]**
- **[edit interfaces *interface-name* gigether-options]**
- **[edit interfaces *interface-name* serial-options]**
- **[edit interfaces *interface-name* sonet-options]**
- **[edit interfaces *interface-name* t1-options]**
- **[edit interfaces *interface-name* t3-options]**

Configuring BERT Testing

To configure BERT:

- Configure the duration of the test.

  [edit interfaces *interface-name interface-type*-options]
  user@host#bert-period *seconds*;

  You can configure the BERT period to last from 1 through 239 seconds on some PICs and from 1 through 240 seconds on other PICs. By default, the BERT period is 10 seconds.

- Configure the error rate to monitor when the inbound pattern is received.

  [edit interfaces *interface-name interface-type*-options]
  user@host#bert-error-rate *rate*;

Copyright © 2018, Juniper Networks, Inc.

rate is the bit error rate. This can be an integer from 0 through 7, which corresponds to a bit error rate from 10^{-0} (1 error per bit) to 10^{-7} (1 error per 10 million bits).

- Configure the bit pattern to send on the transmit path.

  [edit interfaces *interface-name interface-type*-options]
  user@host#bert-algorithm *algorithm*;

  algorithm is the pattern to send in the bit stream. For a list of supported algorithms, enter a **?** after the **bert-algorithm** statement; for example:

```
[edit interfaces t1-0/0/0 t1-options]

user@host# set bert-algorithm ?
Possible completions:
  pseudo-2e11-o152      Pattern is 2^11 -1 (per 0.152 standard)
  pseudo-2e15-o151      Pattern is 2^15 - 1 (per 0.152 standard)
  pseudo-2e20-o151      Pattern is 2^20 - 1 (per 0.151 standard)
  pseudo-2e20-o153      Pattern is 2^20 - 1 (per 0.153 standard)
        ...
```

For specific hierarchy information, see the individual interface types.

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

```
pseudo-2e11-o152      Pattern is 2^11 -1 (per 0.152 standard)
pseudo-2e15-o151      Pattern is 2^15 - 1 (per 0.151 standard)
pseudo-2e20-o151      Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e23-o151      Pattern is 2^23 (per 0.151 standard)
```

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

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

```
all-ones-repeating    Repeating one bits
all-zeros-repeating   Repeating zero bits
alternating-double-ones-zeros  Alternating pairs of ones and zeros
alternating-ones-zeros  Alternating ones and zeros
pseudo-2e11-o152      Pattern is 2^11 -1 (per 0.152 standard)
pseudo-2e15-o151      Pattern is 2^15 - 1 (per 0.151 standard)
pseudo-2e20-o151      Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e7            Pattern is 2^7 - 1
pseudo-2e9-o153       Pattern is 2^9 - 1 (per 0.153 standard)
repeating-1-in-4      1 bit in 4 is set
repeating-1-in-8      1 bit in 8 is set
repeating-3-in-24     3 bits in 24 are set
```

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

 Copyright © 2018, Juniper Networks, Inc.

NOTE: The IQE PICs support only the following algorithms:

```
all-ones-repeating   Repeating one bits
all-zeros-repeating  Repeating zero bits
alternating-double-ones-zeros  Alternating pairs of ones and zeros
alternating-ones-zeros  Alternating ones and zeros
pseudo-2e9-o153      Pattern is 2^9  -1 (per 0.153 (511 type) standard)
pseudo-2e11-o152     Pattern is 2^11 -1 (per 0.152 and 0.153 (2047 type)
 standards)
pseudo-2e15-o151     Pattern is 2^15 -1 (per 0.151 standard)
pseudo-2e20-o151     Pattern is 2^20 -1 (per 0.151 standard)
pseudo-2e20-o153     Pattern is 2^20 -1 (per 0.153 standard)
pseudo-2e23-o151     Pattern is 2^23 -1 (per 0.151 standard)
repeating-1-in-4     1 bit in 4 is set
repeating-1-in-8     1 bit in 8 is set
repeating-3-in-24    3 bits in 24 are set
```

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

NOTE: BERT is supported on the PDH interfaces of the Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP and the DS3/E3 MIC. The following BERT algorithms are supported:

```
all-ones-repeating            Repeating one bits
all-zeros-repeating           Repeating zero bits
alternating-double-ones-zeros Alternating pairs of ones and zeros
alternating-ones-zeros        Alternating ones and zeros
repeating-1-in-4              1 bit in 4 is set
repeating-1-in-8              1 bit in 8 is set
repeating-3-in-24             3 bits in 24 are set
pseudo-2e9-o153               Pattern is 2^9 - 1 (per 0.153 standard)
pseudo-2e11-o152              Pattern is 2^11 - 1 (per 0.152 standard)
pseudo-2e15-o151              Pattern is 2^15 - 1 (per 0.151 standard)
pseudo-2e20-o151              Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e20-o153              Pattern is 2^20 - 1 (per 0.153 standard)
pseudo-2e23-o151              Pattern is 2^23 (per 0.151 standard)
```

Table 52 on page 365 shows the BERT capabilities for various interface types.

Table 52: BERT Capabilities by Interface Type

Interface	T1 BERT	T3 BERT	Comments
12-port T1/E1 Circuit Emulation	Yes (ports 0–11)	—	• Limited algorithms
4-port Channelized OC3/STM1 Circuit Emulation	Yes (port 0–3)	—	• Limited algorithms

Copyright © 2018, Juniper Networks, Inc.

Table 52: BERT Capabilities by Interface Type (continued)

Interface	T1 BERT	T3 BERT	Comments
E1 or T1	Yes (port 0–3)	Yes (port 0–3)	• Single port at a time • Limited algorithms
E3 or T3	Yes (port 0–3)	Yes (port 0–3)	• Single port at a time
Channelized OC12	—	Yes (channel 0–11)	• Single channel at a time • Limited algorithms • No bit count
Channelized STM1	Yes (channel 0–62)	—	• Multiple channels • Only one algorithm • No error insert • No bit count
Channelized T3 and Multichannel T3	Yes (channel 0–27)	Yes (port 0–3 on channel 0)	• Multiple ports and channels • Limited algorithms for T1 • No error insert for T1 • No bit count for T1

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

Starting and Stopping a BERT Test

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

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

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

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

The test runs for the duration you specify with the **bert-period** statement. If you want to terminate the test sooner, issue the **test interface *interface-name* *interface-type*-bert-stop** command:

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

For example:

```
user@host> test interface t3-1/2/0 t3-bert-start
user@host> test interface t3-1/2/0 t3-bert-stop
```

To view the results of the BERT test, issue the **show interfaces extensive | find BERT** command:

Copyright © 2018, Juniper Networks, Inc.

user@host> **show interfaces *interface-name* extensive | find BERT**

For more information about running and evaluating the results of the BERT procedure, see the CLI Explorer.

NOTE: To exchange BERT patterns between a local router and a remote router, include the **loopback remote** statement in the interface configuration at the remote end of the link. From the local router, issue the **test interface** command.

Related Documentation

- *show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port)*

Troubleshooting: em0 Management Interface Link is Down

Problem

Description: Ethernet Link Down alarm is raised when you run the **show chassis alarm** operational mode command on a T640 router, a T1600 router, T4000 router, or a TX Matrix Plus router.

Diagnosis

Perform the following tests to check if the em0 management interface is down on the master Routing Engine or the backup Routing Engine:

1. Run the **show chassis alarms** command.

show chassis alarms

```
user@host0> show chassis alarms
1 alarms currently active
Alarm time Class Description
2011-10-19 11:13:02 MYT Major Host 1 em0 : Ethernet Link Down
```

   Is the alarm **Ethernet Link Down** displayed against the em0 interface of the master Routing Engine (Host 0)?

   - Yes: Contact JTAC for further assistance.
   - No: Continue to the next diagnostic test.

2. Run the **show interfaces em0** and the **show interfaces em0 terse** operational mode commands.

show interfaces em0

```
user@host> show interfaces em0
Physical interface: em0, Enabled, Physical link is Up
Interface index: 1, SNMP ifIndex: 1
Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 100mbps
Device flags : Present Running
Interface flags: SNMP-Traps
...
```

Copyright © 2018, Juniper Networks, Inc.

show interfaces em0 terse

```
user@host> show interfaces em0 terse
Interface Admin Link Proto Local Remote
em0 up up
em0.0 up up inet 10.100.100.1/30
```

Is the em0 interface on the master Routing Engine **up**?

- Yes: Continue to resolution.
- No: Contact JTAC for further assistance

Resolution *To Resolve This Issue*

From the aforementioned diagnosis, we ascertain that the chassis alarm has been raised for the em0 management interface in the backup Routing Engine (Host 1) and not for the master Routing Engine (Host 0).

Implement one of the following solutions on the backup Routing Engine to resolve this issue:

- Disable the em0 interface in the backup Routing Engine:
  1. In configuration mode, go to the **[edit groups re1]** hierarchy level.

     user@host1# **edit groups re1**

  2. Disable the em0 interface.

     [edit groups re1]
     user@host1# **set interfaces em0 disable**

- Ignore the alarm:
  1. In configuration mode, go to the **[edit chassis]** hierarchy level.

     user@host1# **edit chassis**

  2. Ignore the **Ethernet link down** alarm on the management interface by setting the **management-ethernet link-down** alarm option to **ignore**.

     [edit chassis]
     user@host1# **set alarm management-ethernet link-down ignore**

Related Documentation

- Supported Routing Engines by Router on page 17
- *show chassis alarms*

 Copyright © 2018, Juniper Networks, Inc.

Troubleshooting: fxp0 Management Interface Link is Down

Problem

Description: Ethernet Link Down alarm is raised when you run the **show chassis alarm** operational mode command on an M Series router, an MX Series router, a T320 router, a T640 router, a T1600 router, or on a TX Matrix router.

Diagnosis

Perform the following tests to check if the fxp0 interface is down on the master Routing Engine or the backup Routing Engine:

1. Run the **show chassis alarms** command.

show chassis alarms

```
user@host0> show chassis alarms
1 alarms currently active
Alarm time Class Description
2011-10-19 11:13:02 MYT Major Host 1 fxp0 : Ethernet Link Down
```

Is the alarm **Ethernet Link Down** displayed against the fxp0 interface of the master Routing Engine (Host 0)?

- Yes: Contact JTAC for further assistance.
- No: Continue to the next diagnostic test.

2. Run the **show interfaces fxp0** and the **show interfaces fxp0 terse** operational mode commands.

show interfaces fxp0

```
user@host> show interfaces fxp0
Physical interface: fxp0, Enabled, Physical link is Up
Interface index: 1, SNMP ifIndex: 1
Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 100mbps
Device flags : Present Running
Interface flags: SNMP-Traps
...
```

show interfaces fxp0 terse

```
user@host> show interfaces fxp0 terse
Interface Admin Link Proto Local Remote
fxp0 up up
fxp0.0 up up inet 10.100.100.1/30
```

Is the fxp0 interface on the master Routing Engine **up**?

- Yes: Continue to resolution.

Copyright © 2018, Juniper Networks, Inc.

- No: Contact JTAC for further assistance

Resolution ***To Resolve This Issue***

From the diagnosis, we ascertain that the chassis alarm has been raised for the fxp0 management interface in the backup Routing Engine (Host 1) and not for the master Routing Engine (Host 0).

Implement one of the following solutions on the backup Routing Engine to avoid this issue:

- Disable the fxp0 interface in the backup Routing Engine:
  1. In configuration mode, go to the **[edit groups re1]** hierarchy level.

     user@host1# **edit groups re1**

  2. Disable the fxp0 interface.

     [edit groups re1]
     user@host1# **set interfaces fxp0 disable**

- Ignore the alarm:
  1. In configuration mode, go to the **[edit chassis]** hierarchy level.

     user@host1# **edit chassis**

  2. Ignore the **Ethernet link down** alarm on the management interface by setting the **management-ethernet link-down** alarm option to **ignore**.

     [edit chassis]
     user@host1# **set alarm management-ethernet link-down ignore**

Related Documentation

- Supported Routing Engines by Router on page 17
- *show chassis alarms*

Troubleshooting: Faulty Ethernet Physical Interface on an M Series, an MX Series, or a T Series Router

You can follow the basic troubleshooting checklist as explained in the following topics from one through five to troubleshoot an Ethernet physical interface on an M Series, MX Series, or a T Series router.

1. Checking the Cable Connection on page 371
2. Checking the Physical Link Status of the Interface on page 372
3. Checking the Interface Statistics in Detail on page 373

 Copyright © 2018, Juniper Networks, Inc.

Checking the Cable Connection

Problem

Description: Packets are not received or transmitted over the Ethernet physical interface.

Diagnosis

1. Is the correct cable connected to the correct port?
   - Yes: Continue to "Checking the Physical Link Status of the Interface" on page 372.
   - No: See "Resolving Cabling Issue" on page 371.

Resolution

Resolving Cabling Issue

Perform one or more of the following steps to resolve the cabling issue:

1. Connect the cable properly on the local and remote ends without any loose connections.
2. Swap the Ethernet cable for a known good cable if the existing cable is damaged.
3. Connect a single-mode fiber cable to a single-mode interface only and a multimode fiber cable to a multimode interface only. To check fiber optic cable integrity, see "Checking Fiber Optic Cable Integrity" on page 371.
4. Connect the correct small form-factor pluggable transceiver (SFP) on both sides of the cable.

Checking Fiber Optic Cable Integrity

To check the integrity of fiber optic cable with an external cable diagnostic testing tool:

NOTE: **A single-mode fiber cable must be connected to a single-mode interface and a multi-mode fiber cable must be connected to a multi-mode interface.**

1. Measure the received light level at the receiver (R_x) port to see whether the received light level is within the receiver specification of the Ethernet interface.

Copyright © 2018, Juniper Networks, Inc.

2. Measure transmitted light level at the transmitter (T_x) port to see whether the transmitted light level is within the transmitter specification of the Ethernet interface.

Checking the Physical Link Status of the Interface

Problem

Description: Unable to transmit and receive packets on the Ethernet interface even though the cable connection is correct.

Solution

To display the physical link status of the interface, run the **show interface *interface-name* media** operational mode command. For example, on the ge-5/0/1 interface.

```
user@host> show interfaces ge-5/0/1 media
Physical interface: ge-5/0/1, Enabled, Physical link is Up
  Interface index: 317, SNMP ifIndex: 1602
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online, Speed-negotiation: Disabled,
  Auto-MDIX: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 2c:6b:f5:4c:26:73, Hardware address: 2c:6b:f5:4c:26:73
  Last flapped   : 2012-11-30 01:25:37 UTC (03:46:55 ago)
  Input rate     : 880 bps (1 pps)
  Output rate    : 312 bps (0 pps)
  Active alarms  : None
  Active defects : None
  MAC statistics:
   Input bytes: 901296, Input packets: 9799, Output bytes: 976587, Output packets:
 10451
  Filter statistics:
    Filtered packets: 68, Padded packets: 0, Output packet errors: 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
  Interface transmit statistics: Disabled
```

For information about **show interfaces *interface-name* media**, see show interfaces.

Diagnosis

1. Are there any connectivity problems such as input errors and packet loss even though the **Enabled** field displays **Physical link is Up** status and the **Active alarms and Active defect** field displays **None**?

   - Yes: Go to "Checking the Interface Statistics in Detail" on page 373.

Copyright © 2018, Juniper Networks, Inc.

- No: Continue to the next diagnostic test.

2. Does the **Enabled** field display **Physical link is Down** status and the **Active alarms and Active defect** field display **Link**?

   - Yes: The interface is either not connected correctly or is not receiving a valid signal. Go to "Resolving Cabling Issue" on page 371.

   - No: Continue.

Checking the Interface Statistics in Detail

Problem **Description:** The physical interface is not working even though the **Enabled** field displays **Physical link is Up** status and the **Active alarms and Active defect** field displays **None**.

Solution To display the interface statistics in detail, run the **show interface *interface-name* extensive** operational command. For example, on ge-5/0/1 interface.

```
user@host> show interfaces ge-5/0/1 extensive
Physical interface: ge-5/0/1, Enabled, Physical link is Up
  Interface index: 317, SNMP ifIndex: 1602, Generation: 322
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online, Speed-negotiation: Disabled,
  Auto-MDIX: 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:f5:4c:26:73, Hardware address: 2c:6b:f5:4c:26:73
  Last flapped   : 2012-11-30 01:25:37 UTC (04:38:32 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :               806283                    0 bps
   Output bytes  :              1153215                  424 bps
   Input  packets:                10818                    0 pps
   Output packets:                11536                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Label-switched interface (LSI) traffic statistics:
   Input  bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
  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, Policed discards: 233060,
```

Copyright © 2018, Juniper Networks, Inc.

```
L3 incompletes: 0, L2 channel errors: 0,
    L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 11, 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                 3216                 3216                    0

    1 expedited-fo                   0                    0                    0

    2 assured-forw                   0                    0                    0

    3 network-cont                8320                 8320                    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:                      Receive         Transmit
    Total octets                       1007655          1082219
    Total packets                        10886            11536
    Unicast packets                       4350             4184
    Broadcast packets                       32               77
    Multicast packets                     6504             7275
    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                   10886
    Input packet rejects                    68
    Input DA rejects                        68
    Input SA rejects                         0
    Output packet count                                   11536
    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: 5
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority
Limit
```

 Copyright © 2018, Juniper Networks, Inc.

```
                                %             bps     %               usec
    0 best-effort              95       950000000    95                  0       low
none
    3 network-control           5        50000000     5                  0       low
none
  Interface transmit statistics: Disabled
```

For information about **show interfaces *interface-name* detail**, see show interfaces.

Diagnosis

1. Does the **Policed discards**, **L2 channel errors**, **Input DA rejects**, or the **Input SA rejects** field display any errors?

   For information about the errors, seeshow interfaces.

   - Yes: Resolve the errors as needed. Resolving these errors is beyond the scope of this topic.
   - No: Continue with "Performing the Loopback Diagnostic Test" on page 375.

Performing the Loopback Diagnostic Test

Problem

Description: The interface cable is connected correctly and there are no alarms or errors associated with the Ethernet physical interface yet the interface is not working.

Solution

To check whether the Ethernet port or PIC is faulty, you must perform the internal loopback test and hardware loopback test.

To perform a internal loopback diagnostic test on an Ethernet interface, for example on ge-5/0/1 interface:

1. In configuration mode, go to the **[edit interfaces *ge-5/0/1*]** hierarchy level.

   [edit]
   user@host# **edit interface *ge-5/0/1***

2. Set the **gigether-options** option as loopback, commit the configuration and quit configuration mode.

   [edit interfaces *ge-5/0/1*
   user@host# **set gigether-options loopback**
   user@host# **commit**
   user@host# **quit**

3. In operational mode, execute the **show interfaces *ge-5/0/1* media** command.

```
user@host> show interfaces ge-5/0/1 media
Physical interface: ge-5/0/1, Enabled, Physical link is Up
  Interface index: 317, SNMP ifIndex: 1602
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, BPDU Error: None,
```

Copyright © 2018, Juniper Networks, Inc.

```
MAC-REWRITE Error: None, Loopback: Enabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
 Remote fault: Online, Speed-negotiation: Disabled,
  Auto-MDIX: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 2c:6b:f5:4c:26:73, Hardware address: 2c:6b:f5:4c:26:73
  Last flapped   : 2012-11-30 01:25:37 UTC (03:46:55 ago)
  Input rate     : 880 bps (1 pps)
  Output rate    : 312 bps (0 pps)
  Active alarms  : None
  Active defects : None
  MAC statistics:
    Input bytes: 901296, Input packets: 9799, Output bytes: 976587, Output
packets: 10451
  Filter statistics:
    Filtered packets: 68, Padded packets: 0, Output packet errors: 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
  Interface transmit statistics: Disabled
```

NOTE: Delete the **loopback** statement after completing your diagnosis.

Execute one of the following steps for a hardware loopback diagnostic test as needed:

- For an Ethernet PIC with a fiber optic interface—Physically loop the T_x and R_x port and check the status of the physical link with the **show interfaces** ***interface-name*** **media** operational mode command.

- For an Ethernet PIC with an RJ-45 Ethernet interface—Build a loopback plug by crossing pin 1 (T_x +) to pin 3 (R_x +) together and pin 2 (T_x -) and pin 6 (R_x -) together and check the status of the physical link with the **show interfaces** ***interface-name*** **media** operational mode command.

NOTE: For information about loopback testing, see *Performing Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces*.

Diagnosis

1. Does the **Enabled** field display **Physical link is Up** status and the **Active alarms and Active defect** field display **None** when you perform the loopback test?

 Copyright © 2018, Juniper Networks, Inc.

- Yes: Go to the “Checking Other Possibilities” on page 377 section.
- No: Continue to the next diagnostic test.

2. When the Ethernet interface is connected to a remote Ethernet device over multiple patch panels, check to see whether the connection can be looped back at the different patch panels so you can conduct a loopback diagnostic test. Is the loopback diagnostic test successful?

   - Yes: Go to the “Checking Other Possibilities” on page 377 section.
   - No: Contact JTAC for further assistance.

Checking Other Possibilities

Problem **Description:** Loopback diagnostic test is successful but unable to transmit and receive packets on the Ethernet interface.

Solution Use the following commands as needed to troubleshoot an Ethernet interface, for example, an ge-5/0/1 interface:

- Run the **show interfaces *interface-name* terse** operational command to check if the physical interface and logical interfaces are administratively disabled. For example, on ge-5/0/1 interface.

```
user@host> show interfaces ge-5/0/1 terse
     Interface               Admin Link Proto    Local                 Remote
     ge-5/0/1                up    up
     ge-5/0/1.0              up    up   inet     20.1.1.2/24
```

Diagnosis 1. Does the physical interface and its corresponding logical interfaces display **down** in the output of the **show interfaces *interface-name* terse** operational mode command?

   - Yes: Enable the interfaces as shown in “To Enable a Physical Interface” on page 378.
   - No: Continue to the next diagnostic test.

2. Are the **speed**, **duplex**, and **auto-negotiation** fields in the output of **show interfaces *interface-name* extensive** operational mode command correctly set for the interface?

   NOTE: Check if the associated Flexible PIC Concentrator (FPC), Modular Port Concentrator (MPC), or Dense Port Concentrator (DPC) and its Modular Interface Card (MIC) or PIC with its 10-gigabit small form-factor pluggable transceiver (XFP) or SFP supports speed and auto-negotiation settings.

Copyright © 2018, Juniper Networks, Inc.

- Yes: Check *Monitoring Fast Ethernet and Gigabit Ethernet Interfaces* for more troubleshooting tips.
- No: Contact JTAC for further assistance.

To Enable a Physical Interface

To enable a physical interface:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level.

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

2. Check if the interface is administratively disabled by executing the **show** command on the interface. For example on ge-5/0/1 interface.

   user@host# **show ge-5/0/1**

   ```
   disable;
   ```

3. Enable the interface and commit.

   [edit interfaces
   user@host# **delete *interface-name* disable**
   user@host# **commit**

See Also
- show interfaces on page 1119

Time Domain Reflectometry on ACX Series Routers Overview

Time Domain Reflectometry (TDR) is a technology used for diagnosing copper cable states. This technique can be used to determine if cabling is at fault when you cannot establish a link. TDR detects the defects by sending a signal through a cable, and reflecting it from the end of the cable. Open circuits, short circuits, sharp bends and other defects in the cable, reflects the signal back, at different amplitudes, depending on the severity of the defect.

Several factors that result in degraded or low-quality cable plants can cause packet loss, suboptimal connection speed, reduced network efficiency, and complete connection failures. These types of problems can occur because of poor cable construction, identification of pair twists, loose connectors, poor contacts between the points, and stretched or broken pairs of cables. Broadcom transceivers enable you to analyze the condition of the cable plant or topology and identify any problems that have occurred. This functionality is effectively used in the following scenarios:

- Troubleshooting during initial network equipment installation.
- Discovery of failures when network problems occur.

 Copyright © 2018, Juniper Networks, Inc.

- Maintenance of optimally functioning cable plants.
- Fault determination during the testing of network equipment in production cable networks.

TDR supports the following capabilities for examination of cable faults on ACX Series routers:

- Cable status pair (open or short)—When the router operates in Gigabit Ethernet mode, all the four pairs (8 wires) are used. Only Pair-A and Pair-B are required to operate in 10/100BASE-T Ethernet mode. If either of these required pairs is open or short-circuited, the transceiver reports the following faults:
  - Any open wire
  - Wires of a particular pair that are shorted
- Distance to fault per pair—Distance at which an open or a short-circuit is detected in meters. This measurement is also termed as cable length. The transceiver reports the following faults:
  - Cable length when the cable status is normal
  - Distance to fault when the cable status is not normal
- Pair Swap—Swapping of twisted-pairs in straight-through and cross-over cable plants are detected.
- Polarity Swap—Each cable pair carries a differential signal from one end to the other end of the cable. Each wire within the pair is assigned a polarity. The wires in a pair are normally connected in a one-to-one form. This connection enables the transmitter at one end to be connected to the receiver at the other end with same polarity. Sometimes, the wiring within the pair is also swapped. This type of connection is called polarity swap. Broadcom transceivers can detect such swapping and automatically adjust the connection to enable the links to operate normally. However, the transceiver reports polarity swaps that it detects in the cable plant.

On 4-port Gigabit Ethernet and 8-port Gigabit Ethernet MICs with copper SFP transceivers (using BCM54880) and 4-port Gigabit Ethernet, 6-port Gigabit Ethernet, and 8-port Gigabit Ethernet MICs with copper and optical SFP transceivers (using BCM54640E PHY), only 10BASE-T pair polarity is supported. 100BASE-T and 1000BASE-T polarities are not supported.

When the Gigabit Ethernet link cannot be established (for example, if only two pairs are present that are fully functional), TDR in the physical layer (PHY) brings down the link to a 100 MB link, which is called a downshift in the link. The physical layer might require 10-20 seconds for the link to come up if a downgrade in wire speed occurs because it attempts to connect at 1000 MB five times before it falls back to 100BASE-TX.

TDR diagnostics is supported only on copper interfaces and not on fiber interfaces.

Keep the following points in mind when you configure TDR:

- If you connect a port undergoing a TDR test to a Gigabit Ethernet interface that is enabled to automatically detect MDI (Media Dependent Interface) and MDIX (Media Dependent Interface with Crossover) port connections, the TDR result might be invalid.
- If you connect a port undergoing a TDR test to a 100BASE-T copper interface, the unused pairs are reported as faulty because the remote end does not terminate these pairs.
- You must not modify the port configuration while the TDR test is running.
- Because of cable characteristics, you need to run the TDR test multiple times to get accurate results.
- Do not change the port status (such as removing the cable at the near or far end) because such a change can result in inaccurate statistics in the results.
- While measuring the cable length or distance to fault (per pair), sometimes, a few cable length inconsistencies might be observed during a TDR test. Broadcom transceivers have the following cable length limitations:
  - For a properly-terminated good cable, the accuracy of the cable length reported is plus or minus 10 meters.
  - If a pair is open or short-circuited, the far-end termination does not affect the computed result for that pair.
  - The accuracy of the measured cable length, when open and short-circuit conditions are detected, is plus or minus 5 meters.
  - The accuracy of a good pair, when one or more pairs are open or short-circuited, is plus or minus 10 meters.
- Polarity swap detection is supported only in 10BASE-T mode.
- The TDR test does not impact the traffic if the interface operates at 10-Gigabit Ethernet per second of bandwidth, which is the default configuration. However, if the speed of the interface is configured to be other than 10-Gigabit Ethernet, running the TDR test affects the traffic.

  TDR diagnostics might bring the link down and initialize the physical layer (PHY) with default configuration to perform its operation.

  When the TDR validation test is completed, the PHY layer resumes operation in the same manner as before the cable diagnostics test was performed. However, link flaps might be momentarily observed. We recommend that you run the TDR test at a speed of 1 gigabit per second, which is the default configuration, to obtain more accurate results.

TDR is supported on the following interfaces on ACX Series routers:

Copyright © 2018, Juniper Networks, Inc.

- On ACX1000 routers, 4 RJ45 (Cu) ports or 8-port Gigabit Ethernet MICs with small form-factor pluggable (SFP) transceivers and RJ45 connectors.

  On ACX1100 routers, 4-port or 8-port Gigabit Ethernet MICs with SFP transceivers and RJ45 connectors.
- On ACX2000 routers, 8-port Gigabit Ethernet MICs with SFP transceivers and RJ45 connectors.
- On ACX2100 and ACX2200 routers, 4-port Gigabit Ethernet MICs with SFP transceivers and RJ45 connectors.
- On ACX4000 routers, 4-port, 6-port, or 8-port Gigabit Ethernet MICs with SFP transceivers and RJ45 connectors.

You must select the media type as copper for the 1-Gigabit Ethernet interfaces. To specify the media type, include the **media-type** statement with the **copper** option at the **[edit interfaces *interface-name*]** hierarchy level. Media type selection is applicable to ports only in slot 2. When media-type is not set, the port accepts either type of connection. The media type is fiber if a transceiver is installed in the SFP connection. If no transceiver is installed, the media type is copper. The COMBO ports (combination ports) on ACX routers support both the copper and fiber-optic media types. On such ports or interfaces, you must configure the media type as copper to run the TDR test.

You can run the TDR test from operational mode and view the success or failure results of the test. To start a test on a specific interface, issue the **request diagnostics tdr start interface *interface-name*** command. To stop the TDR test currently in progress on the specified interface, issue the **request diagnostics tdr abort interface *interface-name*** command. To display the test results for all copper interfaces, enter the **show diagnostics tdr** command. To display the test results for a particular interface, enter the **show diagnostics tdr interface interface-name** command.

Related Documentation

- Diagnosing a Faulty Twisted-Pair Cable on ACX Series Routers on page 381

Diagnosing a Faulty Twisted-Pair Cable on ACX Series Routers

Problem

Description: A 10/100BASE-T Ethernet interface has connectivity problems that you suspect might be caused by a faulty cable.

Solution

Use the time domain reflectometry (TDR) test to determine whether a twisted-pair Ethernet cable is faulty.

The TDR test:

- Detects and reports faults for each twisted pair in an Ethernet cable. Faults detected include open circuits, short circuits, and impedance mismatches.
- Reports the distance to fault to within 1 meter.
- Detects and reports pair swaps, pair polarity reversals, and excessive pair skew.

Copyright © 2018, Juniper Networks, Inc.

The TDR test is supported on the following ACX routers and interfaces:

- On ACX1000 routers, 4 RJ45 (Cu) ports or 8-port Gigabit Ethernet MICs with small form-factor pluggable (SFP) transceivers and RJ45 connectors.
- On ACX1100 routers, 4-port or 8-port Gigabit Ethernet MICs with SFP transceivers and RJ45 connectors.
- On ACX2000 routers, 8-port Gigabit Ethernet MICs with SFP transceivers and RJ45 connectors.
- On ACX2100 and ACX2200 routers, 4-port Gigabit Ethernet MICs with SFP transceivers and RJ45 connectors.
- On ACX4000 routers, 4-port, 6-port, or 8-port Gigabit Ethernet MICs with SFP transceivers and RJ45 connectors.

NOTE: We recommend running the TDR test on an interface when there is no traffic on the interface.

TDR diagnostics are applicable for copper ports only and not for optical fiber ports.

To diagnose a cable problem by running the TDR test:

1. Run the **request diagnostics tdr** command.

```
user@host> request diagnostics tdr start interface ge-0/0/10

Interface TDR detail:
Test status                     : Test successfully executed  ge-0/0/10
```

2. View the results of the TDR test with the **show diagnostics tdr** command.

```
user@host> show diagnostics tdr interface ge-0/0/10

Interface TDR detail:
Interface name                  : ge-0/0/10
Test status                     : Passed
Link status                     : Down
MDI pair                        : 1-2
  Cable status                  : Normal
  Distance fault                : 0 Meters
  Polartiy swap                 : N/A
  Skew time                     : N/A
MDI pair                        : 3-6
  Cable status                  : Normal
  Distance fault                : 0 Meters
  Polartiy swap                 : N/A
  Skew time                     : N/A
MDI pair                        : 4-5
  Cable status                  : Open
  Distance fault                : 1 Meters
  Polartiy swap                 : N/A
  Skew time                     : N/A
```

 Copyright © 2018, Juniper Networks, Inc.

```
MDI pair                        : 7-8
  Cable status                  : Normal
  Distance fault                : 0 Meters
  Polartiy swap                 : N/A
  Skew time                     : N/A
Channel pair                    : 1
  Pair swap                     : N/A
Channel pair                    : 2
  Pair swap                     : N/A
Downshift                       : N/A
```

3. Examine the **Cable status** field for the four MDI pairs to determine if the cable has a fault. In the preceding example, the twisted pair on pins 4 and 5 is broken or cut at approximately one meter from the **ge-0/0/10** port connection.

NOTE: The **Test Status** field indicates the status of the TDR test, not the cable. The value **Passed** means the test completed—it does not mean that the cable has no faults.

The following is additional information about the TDR test:

- The TDR test can take some seconds to complete. If the test is still running when you execute the **show diagnostics tdr** command, the **Test status** field displays **Started**. For example:

```
user@host> show diagnostics tdr interface ge-0/0/22

Interface TDR detail:
Interface name                 : ge-0/0/22
Test status                    : Started
```

- You can terminate a running TDR test before it completes by using the **request diagnostics tdr abort interface *interface-name*** command. The test terminates with no results, and the results from any previous test are cleared.
- You can display summary information about the last TDR test results for all interfaces on the router that support the TDR test by not specifying an interface name with the **show diagnostics tdr** command. For example:

```
user@host> show diagnostics tdr
Interface    Test status    Link status  Cable status  Max distance fault
 ge-0/0/0         Passed     UP            OK              0
 ge-0/0/1    Not Started    N/A           N/A            N/A
 ge-0/0/2         Passed     UP            OK              0
 ge-0/0/3    Not Started    N/A           N/A            N/A
 ge-0/0/4         Passed     UP            OK              0
 ge-0/0/5         Passed     UP            OK              0
 ge-0/0/6         Passed     UP            OK              0
 ge-0/0/7    Not Started    N/A           N/A            N/A
 ge-0/0/8         Passed    Down           OK              0
 ge-0/0/9    Not Started    N/A           N/A            N/A
ge-0/0/10         Passed    Down           Fault           1
ge-0/0/11         Passed     UP            OK              0
```

Copyright © 2018, Juniper Networks, Inc.

```
ge-0/0/12   Not Started    N/A          N/A           N/A
ge-0/0/13   Not Started    N/A          N/A           N/A
ge-0/0/14   Not Started    N/A          N/A           N/A
ge-0/0/15   Not Started    N/A          N/A           N/A
ge-0/0/16   Not Started    N/A          N/A           N/A
ge-0/0/17   Not Started    N/A          N/A           N/A
ge-0/0/18   Not Started    N/A          N/A           N/A
ge-0/0/19        Passed    Down          OK             0
ge-0/0/20   Not Started    N/A          N/A           N/A
ge-0/0/21   Not Started    N/A          N/A           N/A
ge-0/0/22        Passed     UP           OK             0
ge-0/0/23   Not Started    N/A          N/A           N/A
```

Related Documentation

- Time Domain Reflectometry on ACX Series Routers Overview on page 378
- *request diagnostics tdr*
- *show diagnostics tdr*

Copyright © 2018, Juniper Networks, Inc.

PART 5

Configuration Statements and Operational Commands

 Copyright © 2018, Juniper Networks, Inc.

CHAPTER 13

Configuration Statements

 Copyright © 2018, Juniper Networks, Inc.

 Copyright © 2018, Juniper Networks, Inc.

 Copyright © 2018, Juniper Networks, Inc.

 Copyright © 2018, Juniper Networks, Inc.

 Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

 Copyright © 2018, Juniper Networks, Inc.

 Copyright © 2018, Juniper Networks, Inc.

 Copyright © 2018, Juniper Networks, Inc.

802.3ad

Syntax

```
802.3ad {
  primary | backup;
  ae interface-number ;
  lacp {
    port-priority priority-number;
  }
  link-index index-number
  distribution-list distribution-list-number
}
```

Hierarchy Level

[edit interfaces *interface-name* fastether-options],
[edit interfaces *interface-name* gigether-options]

Release Information

Statement introduced before Junos OS Release 7.4.
primary and **backup** options added in Junos OS Release 8.3.

Description

Specify aggregated Ethernet logical interface number.

Options

bundle—Join an aggregated Ethernet interface.

ae *interface-number*—Aggregated Ethernet logical interface number. For MX Series routers running Junos release 14.2R3 and later you can configure a maximum of 1000 aggregated interfaces. On MX2010 and MX2020 routers you can configure a maximum of 800 aggregated interfaces.

primary | backup—For link protection configurations, specify the link as primary link or backup link for egress traffic.

lacp—Configure Link Aggregation Control Protocol. Specify the port priority in the range 0 through 65535. Default port-priority is 127.

link-index—Specify the desired child link index within the aggregated Ethernet Interface. Index number of the logical interface reflects its initialization sequence.

distribution-list—For targeted distribution, specify the distribution list to which the interface belongs.

Required Privilege Level

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

Related Documentation

- *Configuring an Aggregated Ethernet Interface*
- *Configuring Aggregated Ethernet Link Protection*

 Copyright © 2018, Juniper Networks, Inc.

accept

Syntax accept (any | dhcp-v4 | dhcp-v6 | inet | inet6 | pppoe);

Hierarchy Level [edit interfaces *interface-name* auto-configure stacked-vlan-ranges dynamic-profile *profile-name*],
[edit interfaces *interface-name* auto-configure vlan-ranges dynamic-profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.5.
dhcp-v4 option added in Junos OS Release 10.0.
dhcp-v6, **inet6** and **pppoe** options added in Junos OS Release 10.2.
any option added in Junos OS Release 10.4.

Description Specify the type of VLAN Ethernet packet accepted by an interface that is associated with a VLAN dynamic profile or stacked VLAN dynamic profile.

Options **any**—Any packet type. Specifies that any incoming packets trigger the dynamic creation of a VLAN with properties determined by the auto-configure interface configuration stanza and associated profile attributes. This option is used when configuring wholesaling in a Layer 2 network.

dhcp-v4—IPv4 DHCP packet type. Specifies that incoming IPv4 DHCP discover packets trigger the dynamic creation of a VLAN with properties determined by the auto-configure interface configuration stanza and associated profile attributes

NOTE: The DHCP-specific **mac-address** and **option-82** options are rejected if the **accept** statement is not set to **dhcp-v4**.

dhcp-v6—IPv6 DHCP packet type. Specifies that incoming IPv6 DHCP discover packets trigger the dynamic creation of a VLAN with properties determined by the auto-configure interface configuration stanza and associated profile attributes.

inet—IPv4 Ethernet and ARP packet type.

inet6—IPv6 Ethernet packet type.

pppoe—Point-to-Point Protocol over Ethernet packet type.

NOTE: The **pppoe** VLAN Ethernet packet type option is supported only for MPC/MIC interfaces.

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

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- *Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs*
- *Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs*
- *Configuring VLAN Interfaces for the Layer 2 Wholesale Solution*
- *Configuring Subscriber Packet Types to Trigger VLAN Authentication*

 Copyright © 2018, Juniper Networks, Inc.

accept-source-mac

Syntax

```
accept-source-mac {
  mac-address mac-address {
    policer {
      input cos-policer-name;
      output cos-policer-name;
    }
  }
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Packet Transport Routers.
Statement introduced in Junos OS Release 13.2 for the QFX Series.

Description

For Gigabit Ethernet intelligent queuing (IQ) interfaces only, accept traffic from and to the specified remote media access control (MAC) address.

The **accept-source-mac** statement is equivalent to the **source-address-filter** statement, which is valid for aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only. To allow the interface to receive packets from specific MAC addresses, include the **accept-source-mac** statement.

On untagged Gigabit Ethernet interfaces, you should not configure the **source-address-filter** statement and the **accept-source-mac** statement simultaneously. On tagged Gigabit Ethernet interfaces, you should not configure the **source-address-filter** statement and the **accept-source-mac** statement with an identical MAC address specified in both filters.

The remaining statements are explained separately. See CLI Explorer.

NOTE: The policer statement is not supported on PTX Series Packet Transport Routers.

NOTE: On QFX platforms, if you configure source MAC addresses for an interface using the *static-mac* or *persistent-learning* statements and later configure a different MAC address for the same interface using the **accept-source-mac** statement, the MAC addresses that you previously configured for the interface remain in the ethernet-switching table and can still be used to send packets to the interface.

Copyright © 2018, Juniper Networks, Inc.

Required Privilege Level

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

Related Documentation

- *Configuring MAC Address Filtering*
- *Configuring MAC Address Filtering on PTX Series Packet Transport Routers*
- source-filtering on page 826

Copyright © 2018, Juniper Networks, Inc.

access-concentrator

Syntax

access-concentrator *name*;

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family pppoe],
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit interfaces *interface-name* unit *logical-unit-number* pppoe-options],
[edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* pppoe-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Support at the **[edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]** and **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]** hierarchy levels introduced in Junos OS Release 10.1.
Support at the **[edit ... family pppoe]** hierarchies introduced in Junos OS Release 11.2.

Description

Configure an alternative access concentrator name in the AC-NAME tag in a PPPoE control packet for use with a dynamic PPPoE subscriber interface. If you do not configure the access concentrator name, the AC-NAME tag contains the system name.

NOTE: The **[edit ... family pppoe]** hierarchies are supported only on MX Series routers with MPCs.

Options

name—Name of the access concentrator.

Required Privilege Level

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

Related Documentation

- *Identifying the Access Concentrator*
- *Configuring the PPPoE Family for an Underlying Interface*
- *Configuring Dynamic PPPoE Subscriber Interfaces*
- *PPPoE Overview*

Copyright © 2018, Juniper Networks, Inc.

access-profile

Syntax access-profile *name*;

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

Release Information Statement introduced before Junos OS Release 7.4.
Support for PAP added in Junos OS Release 8.3.
Support for VLAN and stacked VLAN ranges added in Junos OS Release 10.0.

Description For CHAP authentication, the mapping between peer names (or “clients”) and the secrets associated with their respective links. For PAP authentication, the peer's username and password.

For Asynchronous Transfer Mode 2 (ATM2) IQ interfaces only, you can configure a Challenge Handshake Authentication Protocol (CHAP) access profile on the logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:

- **atm-ppp-llc**—PPP over AAL5 logical link control (LLC) encapsulation.
- **atm-ppp-vc-mux**—PPP over AAL5 multiplex encapsulation.

Options *name*—Name of the access profile.

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

Related Documentation

- Configuring the PPP Challenge Handshake Authentication Protocol on page 136
- Configuring the PPP Password Authentication Protocol On a Physical Interface on page 139
- *Junos OS Administration Library*

 Copyright © 2018, Juniper Networks, Inc.

accounting

Syntax

```
accounting {
  destination-class-usage;
  source-class-usage {
    direction;
  }
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Enable IP packet counters on an interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- Enabling Source Class and Destination Class Usage on page 262

accounting-profile

Syntax accounting-profile *name*;

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

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 15.1F6 for PTX Series routers with third-generation FPCs installed.

Description Enable collection of accounting data for the specified physical or logical interface or interface range.

Options ***name***—Name of the accounting profile.

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

Related Documentation

- *Applying an Accounting Profile to the Physical Interface*
- *Applying an Accounting Profile to the Logical Interface*

 Copyright © 2018, Juniper Networks, Inc.

acfc

Syntax
acfc;

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

Release Information
Statement introduced before Junos OS Release 7.4.

Description
For interfaces with PPP encapsulation, configure compression of the Data Link Layer address and control fields. The **acfc** option is not supported with **frame-relay-ppp** encapsulation.

On M320, M120, and T Series routers, address and control field compression (ACFC) is not supported for any ISO family protocols. Do not include the **acfc** statement at the **[edit interfaces *interface-name* ppp-options compression]** hierarchy level when you include the **family iso** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

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

Related Documentation
- Configuring PPP Address and Control Field Compression on page 149

Copyright © 2018, Juniper Networks, Inc.

acknowledge-retries

Syntax acknowledge-retries *number*;

Hierarchy Level [edit interfaces *interface-name* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services and voice services interfaces only, configure the number of retransmission attempts to be made for consecutive hello or remove link messages following the expiration of the acknowledgment timer.

Options *number*—Number of retransmission attempts to be made following the expiration of the acknowledgment timer.

Range: 1 through 5

Default: 2

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*
- action-red-differential-delay on page 421
- hello-timer on page 637

 Copyright © 2018, Juniper Networks, Inc.

acknowledge-timer

Syntax acknowledge-timer *milliseconds*;

Hierarchy Level [edit interfaces *interface-name* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services and voice services interfaces only, configure the maximum time, in milliseconds, to wait for an add link acknowledgment, hello acknowledgment, or remove link acknowledgment message.

Options milliseconds—Time, in milliseconds, to wait for an add link acknowledgment, hello acknowledgment, or remove link acknowledgment message.

Range: 1 through 10 milliseconds

Default: 4 milliseconds

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*
- address on page 424, hello-timer on page 637
- hello-timer on page 637

action (OAM)

Syntax

```
action {
  link-down;
  send-critical-event;
  syslog;
}
```

Hierarchy Level [edit protocols oam ethernet link-fault-management action-profile]

Release Information Statement introduced in Junos OS Release 8.5.

Description Define the action or actions to be taken when the OAM fault event occurs.

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

Related Documentation

- *Specifying the Actions to Be Taken for Link-Fault Management Events*

action (Policer)

Syntax
```
action {
  loss-priority high then discard;
}
```

Hierarchy Level [edit firewall three-color-policer *policer-name*]

Release Information Statement introduced in Junos OS Release 8.2.

Description This statement discards high loss priority traffic as part of a configuration using tricolor marking on a logical interface.

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

Related Documentation
- *Class of Service Feature Guide for Routing Devices and EX9200 Switches*
- logical-interface-policer on page 747

action-profile (Applying to CFM)

Syntax action-profile *profile-name*;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id* remote-mep *mep-id*]

Release Information Statement introduced in Junos OS Release 8.4.

Description Identify the action profile to use.

Options ***profile-name***—Name of the action profile to use.

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

Related Documentation
- *Configuring a MEP to Generate and Respond to CFM Protocol Messages*

Copyright © 2018, Juniper Networks, Inc.

action-profile (Defining for CFM)

Syntax

```
action-profile profile-name {
  event {
    ais-trigger-condition {
      adjacency-loss;
      all-defects;
      cross-connect-ccm;
      erroneous-ccm;
      receive-ais;
    }
    interface-status-tlv (down | lower-layer-down);
    port-status-tlv blocked;
    rdi;
  }
  action {
    interface-down;
    log-and-generate-ais {
      interval(1m | 1s);
      levelvalue;
      priority value;
    }
  }
  default-actions {
    interface-down;
  }
}
```

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management]

Release Information Statement introduced in Junos OS Release 8.4.

Description Configure a name and default action for an action profile.

Options *profile-name*—Name of the action profile.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *Configuring a CFM Action Profile to Specify CFM Actions for CFM Events*
- default-actions on page 516
- *event (CFM)*
- interface-down on page 686

action-profile

List of Syntax

Syntax: T, M, MX and ACX Series Routers, SRX Series Firewalls and EX Series Switches on page 420
Syntax: EX Series Switches and NFX Series Devices on page 420

Syntax: T, M, MX and ACX Series Routers, SRX Series Firewalls and EX Series Switches

```
action-profile profile-name {
  action {
    link-down;
    send-critical-event;
    syslog;
  }
  event {
    link-adjacency-loss;
    link-event-rate {
      frame-error count;
      frame-period count;
      frame-period-summary count;
      symbol-period count;
    }
    protocol-down;
  }
}
```

Syntax: EX Series Switches and NFX Series Devices

```
action-profile profile-name;
  action {
    syslog;
    link-down;
  }
  event {
    link-adjacency-loss;
    link-event-rate {
      frame-error  count;
      frame-period count;
      frame-period-summary  count;
      symbol-period count;
    }
  }
```

Hierarchy Level

[edit protocols oam ethernet link-fault-management]

Release Information

Statement introduced in Junos OS Release 8.5 for T, M, MX and ACX Series Routers, SRX Series Firewalls, and EX Series Switches, .
Statement introduced in Junos OS Release 9.4 for EX Series switches and NFX Series devices.

Description

Configure an Ethernet OAM link fault management (LFM) action profile by specifying a profile name.

The remaining statements are explained separately. See CLI Explorer.

Copyright © 2018, Juniper Networks, Inc.

Options ***profile-name***—Name of the action profile.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *Configuring an OAM Action Profile*
- *Configuring Ethernet OAM Link Fault Management (CLI Procedure)*

action-red-differential-delay

Syntax action-red-differential-delay (disable-tx | remove-link);

Hierarchy Level [edit interfaces *interface-name* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services and voice services interfaces only, configure the action to be taken when the differential delay exceeds the red limit.

Options **disable-tx**—Disable transmission on the bundle link.

remove-link—Remove bundle link from service.

Default: disable-tx

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*
- remote on page 926
- yellow-differential-delay on page 1106

Copyright © 2018, Juniper Networks, Inc.

activation-delay

Syntax activation-delay *seconds*;

Hierarchy Level [edit interfaces dl*n* unit *logical-unit-number* dialer-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description (J Series Services Routers) For ISDN interfaces, configure the ISDN dialer activation delay. Used only for dialer backup and dialer watch cases.

Options *seconds*—Interval before the backup interface is activated after the primary interface has gone down.

Range: 1 through 4,294,967,295 seconds

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

Related Documentation

- *Junos OS Interfaces and Routing Configuration Guide*

Copyright © 2018, Juniper Networks, Inc.

activation-priority

Syntax

activation-priority *priority*;

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* dynamic-call-admission-control],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* dynamic-call-admission-control]

Release Information

Statement introduced in Junos OS Release 8.2.

Description

(J4350 and J6350 Services Routers supporting voice over IP with the TGM550 media gateway module) For Fast Ethernet and Gigabit Ethernet interfaces, ISDN BRI interfaces, and serial interfaces with PPP or Frame Relay encapsulation, configure the dynamic call admission control (dynamic CAC) activation priority value.

Options

priority—The activation priority in which the interface is used for providing call bandwidth. The interface with the highest activation priority value is used as the primary link for providing call bandwidth. If the primary link becomes unavailable, the TGM550 switches over to the next active interface with the highest activation priority value, and so on.

Range: 0 through 255

Default: 50

Required Privilege Level

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*
- *Junos OS Interfaces and Routing Configuration Guide*

Copyright © 2018, Juniper Networks, Inc.

address

Syntax

```
address address {
  arp ip-address (mac | multicast-mac) mac-address <publish>;
  broadcast address;
  destination address;
  destination-profile name;
  eui-64;
  master-only;
  multipoint-destination address dlci dlci-identifier;
  multipoint-destination address {
    epd-threshold cells;
    inverse-arp;
    oam-liveness {
      up-count cells;
      down-count cells;
    }
    oam-period (disable | seconds);
    shaping {
      (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate burst
        length);
      queue-length number;
    }
    vci vpi-identifier.vci-identifier;
  }
  primary;
  preferred;
  virtual-gateway-address
  (vrrp-group | vrrp-inet6-group) group-number {
    (accept-data | no-accept-data);
    advertise–interval seconds;
    authentication-type authentication;
    authentication-key key;
    fast-interval milliseconds;
    (preempt | no-preempt) {
      hold-time seconds;
    }
    priority-number number;
    track {
      priority-cost seconds;
      priority-hold-time interface-name {
        interface priority;
        bandwidth-threshold bits-per-second {
          priority;
        }
      }
      route ip-address/mask routing-instance instance-name priority-cost cost;
    }
    virtual-address [ addresses ];
  }
}
```

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

 Copyright © 2018, Juniper Networks, Inc.

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

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure the interface address.

Options *address*—Address of the interface.

- In Junos OS Release 13.3 and later, when you configure an IPv6 host address and an IPv6 subnet address on an interface, the commit operation fails.
- In releases earlier than Junos OS Release 13.3, when you use the same configuration on an interface, the commit operation succeeds, but only one of the IPv6 addresses that was entered is assigned to the interface. The other address is not applied.

NOTE: If you configure the same address on multiple interfaces in the same routing instance, Junos OS uses only the first configuration, and the remaining address configurations are ignored and can leave interfaces without an address. Interfaces that do not have an assigned address cannot be used as a donor interface for an unnumbered Ethernet interface.

For example, in the following configuration the address configuration of interface xe-0/0/1.0 is ignored:

```
interfaces {
  xe-0/0/0 {
    unit 0 {
      family inet {
        address 192.168.1.1/8;
      }
    }
  }
  xe-0/0/1 {
    unit 0 {
     family inet {
       address 192.168.1.1/8;
     }
  }
}
```

For more information on configuring the same address on multiple interfaces, see "Configuring the Interface Address" on page 207.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

NOTE: The edit logical-systems hierarchy is not available on QFabric systems.

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

Related Documentation

- Configuring the Protocol Family on page 206
- *Junos OS Administration Library*
- *family*
- negotiate-address on page 805
- unnumbered-address (Ethernet) on page 1066

advertise-interval

Syntax advertise-interval *milliseconds*;

Hierarchy Level [edit interfaces *interface-name* sonet-options aps]

Release Information Statement introduced before Junos OS Release 7.4.

Description Modify the Automatic Protection Switching (APS) interval at which the protect and working routers send packets to their neighbors to advertise that they are operational. A router considers its neighbor to be operational for a period, called the hold time, that is, by default, three times the advertisement interval.

Options *milliseconds*—Interval between advertisement packets.

Range: 1 through 65,534 milliseconds

Default: 1000 milliseconds

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

Related Documentation

- *Configuring APS Timers*

 Copyright © 2018, Juniper Networks, Inc.

age

Syntax age (30m | 10m | 1m | 30s | 10s);

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management linktrace]

Release Information Statement introduced in Junos OS Release 8.5.

Description Time to wait (in minutes or seconds) for a response. If no response is received, the request and response entry is deleted from the linktrace database.

Default 10 minutes

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

Related Documentation

- *Configuring Linktrace Protocol in CFM*

Copyright © 2018, Juniper Networks, Inc.

agent-specifier

Syntax

```
agent-specifier {
  aci circuit-id-string ari remote-id-string {
    drop;
    delay seconds;
    terminate;
    dynamic-profile profile-name;
    routing-instance routing-instance-name;
    static-interface interface-name;
  }
}
```

Hierarchy Level [edit protocols pppoe service-name-tables *table-name* service *service-name*]

Release Information Statement introduced in Junos OS Release 10.0.

drop, **delay**, **terminate**, **dynamic-profile**, **routing-instance**, and **static-interface** options introduced in Junos OS Release 10.2.

Description Specify the action taken by the interface for the specified agent circuit identifier/agent remote identifier (ACI/ARI) pair when the interface receives a PPPoE Active Discovery Initiation (PADI) control packet that includes the vendor-specific tag with ACI/ARI pair information. You can configure an ACI/ARI pair for a named service, **empty** service, or **any** service in a PPPoE service name table. A maximum of 8000 ACI/ARI pairs are supported per PPPoE service name table. You can distribute the ACI/ARI pairs in any combination among the named, **empty**, and **any** service entries in the service name table.

You can use an asterisk (*) as a wildcard character to match ACI/ARI pairs, the ACI alone, or the ARI alone. The asterisk can be placed only at the beginning, the end, or both the beginning and end of the identifier string. You can also specify an asterisk alone for either the ACI or the ARI. You cannot specify only an asterisk for both the ACI and the ARI. When you specify a single asterisk as the identifier, that identifier is ignored in the PADI packet.

For example, suppose you care about matching only the ACI and do not care what value the ARI has in the PADI packet, or even whether the packet contains an ARI value. In this case you can set the ***remote-id-string*** to a single asterisk. Then the interface ignores the ARI received in the packet and the interface takes action based only on matching the specified ACI.

Default The default action is terminate.

Options **aci *circuit-id-string***—Identifier for the agent circuit ID that corresponds to the DSLAM interface that initiated the service request. This is a string of up to 63 characters.

ari *remote-id-string*—Identifier for the subscriber associated with the DSLAM interface that initiated the service request. This is a string of up to 63 characters.

 Copyright © 2018, Juniper Networks, Inc.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

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

Related Documentation
- *Configuring PPPoE Service Name Tables*
- *Assigning an ACI/ARI Pair to a Service Name and Configuring the Action Taken When the Client Request Includes ACI/ARI Information*

aggregate (Gigabit Ethernet CoS Policer)

Syntax

```
aggregate {
  bandwidth-limit bps;
  burst-size-limit bytes;
}
```

Hierarchy Level
[edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile policer *cos-policer-name*]

Release Information
Statement introduced before Junos OS Release 7.4.

Description
Define a policer to apply to nonpremium traffic.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation
- *Configuring Gigabit Ethernet Policers*
- premium (Hierarchical Policer) on page 898
- ieee802.1p on page 652

Copyright © 2018, Juniper Networks, Inc.

aggregate (Hierarchical Policer)

Syntax

```
aggregate {
  if-exceeding {
    bandwidth-limit bandwidth;
    burst-size-limit burst;
  }
  then {
    discard;
  }
}
```

Hierarchy Level [edit firewall hierarchical-policer]

Release Information Statement introduced in Junos OS Release 9.5.

Description On M40e, M120, and M320 (with FFPC and SFPC) edge routers and T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs, T4000 routers with Type 5 FPC and Enhanced Scaling Type 4 FPC, configure an aggregate hierarchical policer.

Options Options are described separately.

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

Related Documentation

- Applying Policers on page 226
- *Class of Service Feature Guide for Routing Devices and EX9200 Switches*

 Copyright © 2018, Juniper Networks, Inc.

aggregate (SONET/SDH)

Syntax aggregate as*x*;

Hierarchy Level [edit interfaces *interface-name* sonet-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify aggregated SONET/SDH logical interface number.

Options **as*x***—Aggregated SONET/SDH logical interface number.

Range: 0 through 15

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

Related Documentation
- *Configuring Aggregated SONET/SDH Interfaces*

aggregate-ports

Syntax aggregate-ports;

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number*]

Release Information Statement introduced in Junos OS Release 8.1.

Description For T Series routers only, specify OC768-over-OC192 mode on the 4-port OC192C PIC. Four OC192 links are aggregated into one OC768 link with one logical interface.

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

Related Documentation
- *Configuring 4-Port OC192 PIC to Operate in OC768-over-OC192 Mode*

Copyright © 2018, Juniper Networks, Inc.

aggregated-ether-options

Syntax

```
aggregated-ether-options {
  ethernet-switch-profile {
    ethernet-policer-profile {
      input-priority-map {
        ieee802.1p premium [ values ];
      }
      output-priority-map {
        classifier {
          premium {
            forwarding-class class-name {
              loss-priority (high | low);
            }
          }
        }
      }
      policer cos-policer-name {
        aggregate {
          bandwidth-limit bps;
          burst-size-limit bytes;
        }
        premium {
          bandwidth-limit bps;
          burst-size-limit bytes;
        }
      }
    }
    (mac-learn-enable | no-mac-learn-enable);
  }
  (flow-control | no-flow-control);
  lacp {
    (active | passive);
    link-protection {
      disable;
    (revertive | non-revertive);
    periodic interval;
    system-priority priority;
    system-id system-id;
  }
  link-protection;
  load-balance;
  link-speed speed;
  logical-interface-chassis-redundancy;
  logical-interface-fpc-redundancy;
  (loopback | no-loopback);
  minimum-links number;
  rebalance-periodic time hour:minute <interval hours>;
  source-address-filter {
    mac-address;
    (source-filtering | no-source-filtering);
  }
}
```


Copyright © 2018, Juniper Networks, Inc.

Hierarchy Level [edit interfaces ae*x*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure aggregated Ethernet-specific interface properties.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *Ethernet Interfaces Overview*

aggregated-sonet-options

Syntax

```
aggregated-sonet-options {
  link-speed speed;
  minimum-links number;
}
```

Hierarchy Level [edit interfaces as*x*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure aggregated SONET/SDH-specific interface properties.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *Configuring Aggregated SONET/SDH Interfaces*

Copyright © 2018, Juniper Networks, Inc.

alarm (optics-options)

Syntax

```
alarm low-light–alarm {
  (link-down | syslog);
}
```

Hierarchy Level [edit interfaces *interface-name* optics-options]

Release Information Statement introduced in Junos OS Release 10.0.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description Specify the action to take if the receiving optics signal is below the optics low-light alarm threshold.

Options **link-down**—Drop the 10-Gigabit Ethernet link and marks link as down.

syslog—Write the optics information to the system log.

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

Related Documentation

- *Configuring 10-Gigabit Ethernet Link Down Notification for Optics Options Alarm or Warning*
- *100-Gigabit Ethernet OTN Options Configuration Overview*

 Copyright © 2018, Juniper Networks, Inc.

alias (Interfaces)

Syntax alias *alias-name*;

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

Release Information Statement introduced in Junos OS Release 13.3.

Description Configure a textual description of a physical interface or the logical unit of an interface to be the alias of an interface name. The alias name can be a single line of text. If the text contains spaces, enclose it in quotation marks. If you configure an alias name, the alias name is displayed instead of the interface name in the output of all **show**, **show interfaces**, and other operational mode commands. In Junos OS Release 12.3R8 and later, display of the alias can be suppressed in favor of the actual interface name by using the **display no-interface-alias** parameter along with the show command.

Options ***alias-name***—Text to denote an easily identifiable, meaningful alias name for the interface. If the text includes spaces, enclose the entire text in quotation marks.

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

Related Documentation

- Example: Adding an Interface Alias Name on page 123
- *Junos OS Network Interfaces Library for Routing Devices*

allow-any-vci

Syntax allow-any-vci;

Hierarchy Level [edit interfaces *interface-name* unit 0],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit 0]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro routers.

Description Dedicate entire ATM device to ATM cell relay circuit.

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

Related Documentation

- *Configuring an ATM1 Cell-Relay Circuit Overview*

allow-fragmentation

Syntax allow-fragmentation;

Hierarchy Level [edit interfaces gr-*fpc/pic/port* unit *logical-unit-number* tunnel],
[edit logical-systems *logical-system-name* interfaces gr-*fpc/pic/port* unit *logical-unit-number* tunnel]

Release Information Statement introduced in Junos OS Release 9.2.
Statement introduced in Junos OS Release 15.1X53-D10 for QFX10000 switches.

Description For a generic routing encapsulation (GRE) tunnel, enable fragmentation of GRE-encapsulated packets whose size exceeds the maximum transmission unit (MTU) value of a link that the packet passes through. The don't-fragment (DF) bit is not set in the outer IP header of GRE-encapsulated packets.

To enable the reassembly of fragmented GRE-encapsulated packets on GRE tunnel interfaces at the endpoint of the GRE tunnel, include the *reassemble-packets* statement for the interface.

Default If you do not include the **allow-fragmentation** statement, fragmentation of GRE-encapsulated packets is disabled.

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

Related Documentation

- *reassemble-packets*
- *Enabling Fragmentation and Reassembly on Packets After GRE-Encapsulation*
- *Junos OS Services Interfaces Library for Routing Devices*

 Copyright © 2018, Juniper Networks, Inc.

allow-remote-loopback

Syntax allow-remote-loopback;

Hierarchy Level [edit protocols oam link-fault-management interface *interface-name* negotiation-options]

Release Information Statement introduced in Junos OS Release 8.4.

Description Enable the remote loopback on IQ2 and IQ2-E Gigabit Ethernet interfaces, and Ethernet interfaces on the MX Series routers and EX Series switches.

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

Related Documentation

- *Enabling Remote Loopback Support on the Local Interface*

annex

Syntax annex (annex-a | annex-b);

Hierarchy Level [edit interfaces *interface-name* shdsl-options],
[edit interfaces *interface-name* sonet-options aps],
[edit logical-systems *logical-system-name* interfaces *interface-name* shdsl-options]

Release Information Statement introduced in Junos OS Release 7.4.

Description For M320 and M120 routers only, for Multiplex Section Protection (MSP) switching on SDH interfaces, set **annex-b**. You must also configure the **working protection circuit** under the **[edit interfaces so-*fpc/pic/port* sonet-options aps]** hierarchy level.

Default **annex-b**

Options **annex-a**—Use for North American SHDSL network implementations.

annex-b—Use for European SHDSL network implementations.

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

apply-action-profile

Syntax apply-action-profile *profile-name*;

Hierarchy Level [edit protocols oam ethernet link-fault-management interface]

Release Information Statement introduced in Junos OS Release 8.5.

Description Apply the specified action profile to the interface for link-fault management.

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

Related Documentation
- *Applying an Action Profile*

 Copyright © 2018, Juniper Networks, Inc.

aps

Syntax

```
aps {
  advertise-interval milliseconds;
  annex-b
  authentication-key key;
  (break-before-make | no-break-before-make);
  fast-aps-switch;
  force;
  hold-time milliseconds;
  lockout;
  neighbor address;
  paired-group group-name;
  preserve-interface;
  protect-circuit group-name;
  request;
  revert-time seconds;
  switching-mode (bidirectional | unidirectional);
  working-circuit group-name;
}
```

Hierarchy Level

[edit interfaces *interface-name* sonet-options]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure Automatic Protection Switching (APS) on the router.

For DS3 channels on a channelized OC12 interface, configure APS on channel 0 only. If you configure APS on channels 1 through 11, it is ignored.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Automatic Protection Switching and Multiplex Section Protection Overview*

Copyright © 2018, Juniper Networks, Inc.

arp (Interfaces)

Syntax

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

```
arp {
  aging-timer minutes;
  gratuitous-arp-delayseconds;
  gratuitous-arp-on-ifup;
  interfaces {
    interface-name {
      aging-timer minutes;
    }
  }
  passive-learning;
  purging;
}
```

Syntax (EX Series)

```
arp {
  aging-timer minutes;
}
```

Hierarchy Level

[edit system]

[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*],

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

NOTE: The **edit logical-systems** hierarchy is not available on QFabric systems.

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

For Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only, configure Address Resolution Protocol (ARP) table entries, mapping IP addresses to MAC addresses. You can enable backup VRRP routers to learn ARP requests for VRRP-IP to VRRP-MAC address translation. You can also set the time interval between ARP updates.

NOTE: By default, an ARP policer is installed that is shared among all the Ethernet interfaces on which you have configured the **family inet** statement. By including the **arp** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet policer]** hierarchy level, you can apply a specific ARP-packet policer to an interface. This feature is not available on EX Series switches.

 Copyright © 2018, Juniper Networks, Inc.

When you need to conserve IP addresses, you can configure an Ethernet interface to be unnumbered by including the **unnumbered-address** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet]** hierarchy level.

NOTE: For EX-Series switches, set only the time interval between ARP updates.

Options **ip-address**—IP address to map to the MAC address. The IP address specified must be part of the subnet defined in the enclosing **address** statement.

mac *mac-address*—MAC address to map to the IP address. Specify the MAC address as six hexadecimal bytes in one of the following formats: ***nnnn.nnnn.nnnn*** or ***nn*:*nn*:*nn*:*nn*:*nn*:*nn***. For example, **0000.5e00.5355** or **00:00:5e:00:53:55**.

multicast-mac *mac-address*—Multicast MAC address to map to the IP address. Specify the multicast MAC address as six hexadecimal bytes in one of the following formats: ***nnnn.nnnn.nnnn*** or ***nn*:*nn*:*nn*:*nn*:*nn*:*nn***. For example, **0000.5e00.5355** or **00:00:5e:00:53:55**.

publish—(Optional) Have the router or switch reply to ARP requests for the specified IP address. If you omit this option, the router or switch uses the entry to reach the destination but does not reply to ARP requests.

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

aging-timer—Time interval in minutes between ARP updates. In environments where the number of ARP entries to update is high (for example, on routers only, metro Ethernet environments), increasing the time between updates can improve system performance.

passive-learning (QFX-Series only)—Configure backup VRRP routers or switches to learn the ARP mappings (IP-to-MAC address) for hosts sending the requests. By default, the backup VRRP router drops these requests; therefore, if the master router fails, the backup router must learn all entries present in the ARP cache of the master router. Configuring passive learning reduces transition delay when the backup router is activated.

Copyright © 2018, Juniper Networks, Inc.

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

Related Documentation

- *Configuring Static ARP Table Entries For Mapping IP Addresses to MAC Addresses*
- *Configuring Junos OS ARP Learning and Aging Options for Mapping IPv4 Network Addresses to MAC Addresses*
- *Junos OS Network Interfaces Library for Routing Devices*
- *Junos OS System Basics Configuration Guide* .

asynchronous-notification

Syntax (asynchronous-notification | no-asynchronous-notification);

Hierarchy Level [edit interfaces ge-*fpc/pic/port* gigether-options]

Release Information Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description (MX Series routers, T Series routers) For all Gigabit Ethernet interfaces (1-Gigabit, 10-Gigabit, and 100-Gigabit), configure support for notification of link down alarm generation and transfer.

(M120 and M320 routers) For all 10-Gigabit Ethernet PIC interfaces, configure support for notification of link down alarm generation and transfer.

- **asynchronous-notification**—Support notification of link down alarm generation and transfer.
- **no-asynchronous-notification**—Prohibit notification of link down alarm generation and transfer.

Default Support for notification of link down alarm generation and transfer is not enabled.

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

Related Documentation

- *Gigabit Ethernet Notification of Link Down Alarm Overview*
- *Configuring Gigabit Ethernet Notification of Link Down Alarm*

 Copyright © 2018, Juniper Networks, Inc.

atm-encapsulation

Syntax atm-encapsulation (direct | plcp);

Hierarchy Level [edit interfaces at-*fpc/pic/port* e3-options],
[edit interfaces at-*fpc/pic/port* t3-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure encapsulation for E3 and T3 traffic over ATM interfaces.

Default Physical Layer Convergence Protocol (PLCP) encapsulation is the default for T3 traffic and for E3 traffic using G.751 framing.

Options **direct**—Use direct encapsulation. G.832 framing on E3 interfaces requires direct encapsulation.

plcp—Use PLCP encapsulation.

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

Related Documentation

- *Configuring E3 and T3 Parameters on ATM Interfaces*
- encapsulation on page 577

Copyright © 2018, Juniper Networks, Inc.

atm-options

Syntax

```
atm-options {
  cell-bundle-size cells;
  ilmi;
  linear-red-profiles profile-name {
    high-plp-max-threshold percent;
    low-plp-max-threshold percent;
    queue-depth cells high-plp-threshold percent low-plp-threshold percent;
  }
  mpls {
    pop-all-labels {
      required-depth number;
    }
  }
  pic-type (atm1 | atm2);
  plp-to-clp;
  promiscuous-mode {
    vpi vpi-identifier;
  }
  scheduler-maps map-name {
    forwarding-class class-name {
      epd-threshold cells plp1 cells;
      linear-red-profile profile-name;
      priority (high | low);
      transmit-weight (cells number | percent number);
    }
    vc-cos-mode (alternate | strict);
  }
  use-null-cw;
  vpi vpi-identifier {
    maximum-vcs maximum-vcs;
    oam-liveness {
      up-count cells;
      down-count cells;
    }
    oam-period (disable | seconds);
    shaping {
      (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate burst
        length);
      queue-length number;
    }
  }
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers.

Description Configure ATM-specific physical interface properties.

 Copyright © 2018, Juniper Networks, Inc.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

NOTE: Certain options apply only to specific platforms.

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

Related Documentation

- Interface Encapsulations Overview on page 46
- multipoint-destination on page 794
- shaping on page 965
- vci on page 1073

atm-scheduler-map

Syntax atm-scheduler-map (*map-name* | default);

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Associate a scheduler map with a virtual circuit on a logical interface.

Options *map-name*—Name of scheduler map that you define at the **[edit interfaces** ***interface-name*** **atm-options scheduler-maps]** hierarchy level.

default—The default scheduler mapping.

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

Related Documentation

- *ATM2 IQ VC Tunnel CoS Components Overview*
- scheduler-maps (For ATM2 IQ Interfaces) on page 950

Copyright © 2018, Juniper Networks, Inc.

authentication

Syntax

```
authentication {
  packet-types [packet-types];
  password password-string;
  username-include {
    circuit-id;
    circuit-type;
    delimiter delimiter-character;
    domain-name domain-name-string;
    interface-name;
    mac-address;
    option-18;
    option-37;
    option-82 <circuit-id> <remote-id>;
    radius-realm radius-realm-string;
    remote-id;
    user-prefix user-prefix-string;
  }
}
```

Hierarchy Level

[edit interfaces *interface-name* auto-configure vlan-ranges],
[edit interfaces *interface-name* auto-configure stacked-vlan-ranges]

Release Information

Statement introduced in Junos OS Release 10.0.

Description

Specify the authentication parameters that trigger the Access-Request message to AAA for the interface.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

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

Related Documentation

- *Subscribers over Static Interfaces Configuration Overview*
- *Configuring the Static Subscriber Global Authentication Password*
- *Configuring a Username for Authentication of Out-of-Band Triggered Dynamic VLANs*
- *Layer 2 Wholesale with ANCP-Triggered VLANs Overview*

Copyright © 2018, Juniper Networks, Inc.

authentication-key

Syntax authentication-key *key*;

Hierarchy Level [edit interfaces *interface-name* sonet-options aps]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the Automatic Protection Switching (APS) authentication key (password).

Options *key*—Authentication password. It can be 1 through 8 characters long. Configure the same key for both the working and protect routers.

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

Related Documentation

- *Configuring Basic Automatic Protect Switching*
- For information about the **authentication-key** statement at the **[edit interfaces *interface-name* unit *unit-number* family inet address *address* (vrrp-group | vrrp-inet6-group) *group-number*]** or **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *unit-number* family (inet | inet6) address *address* (vrrp-group | vrrp-inet6-group) *group-number*]** hierarchy level, see the *Junos OS High Availability Library for Routing Devices*.

authentication-profile-name

Syntax authentication-profile-name *access-profile-name*;

Hierarchy Level [edit protocols dot1x authenticator]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the RADIUS authentication profile to use for user authentication when establishing an IEEE 802.1x Port-Based Network Access Control (**dot1x**) connection.

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

Related Documentation

- *IEEE 802.1x Port-Based Network Access Control Overview*
- authenticator on page 448
- dot1x on page 543

authenticator

Syntax

```
authenticator {
  authentication-profile-name access-profile-name;
  interface interface-id {
    maximum-requests integer;
    quiet-period seconds;
    reauthentication (disable | interval seconds);
    retries integer;
    server-timeout seconds;
    supplicant (single);
    supplicant-timeout seconds;
    transmit-period seconds;
  }
}
```

Hierarchy Level [edit protocols dot1x]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify an authentication profile for user or client authentication and configure the Ethernet interface for 802.1x protocol operation.

Options **authentication-profile-name *access-profile-name***—Specifies the RADIUS authentication profile for user or client authentication.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *IEEE 802.1x Port-Based Network Access Control Overview*
- authentication-profile-name on page 447
- dot1x on page 543

 Copyright © 2018, Juniper Networks, Inc.

auto-configure

Syntax

```
auto-configure {
  vlan-ranges {
    access-profile profile-name;
    authentication {
      packet-types [packet-types];
      password password-string;
      username-include{
        circuit-id;
        circuit-type;
        delimiter delimiter-character;
        domain-name domain-name-string;
        interface-name;
        mac-address;
        option-18;
        option-37;
        option-82 <circuit-id> <remote-id>;
        radius-realm radius-realm-string;
        remote-id;
        user-prefix user-prefix-string;
      }
    }
    dynamic-profile profile-name {
      accept (any | dhcp-v4 | dhcp-v6 | inet | inet6 | pppoe);
      accept-out-of-band protocol;
      ranges (any | low-tag)–(any | high-tag);
    }
    override;
  }
  stacked-vlan-ranges {
    access-profile profile-name;
    authentication {
      packet-types [packet-types];
      password password-string;
      username-include {
        circuit-type;
        delimiter delimiter-character;
        domain-name domain-name-string;
        interface-name;
        mac-address;
        option-18;
        option-37;
        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 | pppoe);
      ranges (any | low-tag–high-tag),(any | low-tag–high-tag);
    }
    override;
  }
```

```
  remove-when-no-subscribers;
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Enable the configuration of dynamic, auto-sensed VLANs.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

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

Related Documentation

- *Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs*
- *Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs*

auto-discovery

Syntax auto-discovery;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id*]

Release Information Statement introduced in Junos OS Release 8.4.

Description Enable the MEP to accept continuity check messages from all remote MEPs.

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

Related Documentation

- *Configuring a MEP to Generate and Respond to CFM Protocol Messages*

 Copyright © 2018, Juniper Networks, Inc.

auto-negotiation

Syntax (auto-negotiation | no-auto-negotiation) <remote-fault (local-interface-online | local-interface-offline)>;

Hierarchy Level [edit interfaces *interface-name* ether-options],
[edit interfaces *interface-name* gigether-options],
[edit interfaces ge-*pim*/0/0 switch-options switch-port *port-number*]

Release Information Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description For Gigabit Ethernet interfaces on M Series, MX Series, T Series, TX Matrix routers, and ACX Series routers explicitly enable autonegotiation and remote fault. For EX Series switches, explicitly enable autonegotiation only.

- **auto-negotiation**—Enables autonegotiation. This is the default.
- **no-auto-negotiation**—Disable autonegotiation. When autonegotiation is disabled, you must explicitly configure the link mode and speed.

When you configure Tri-Rate Ethernet copper interfaces to operate at 1 Gbps, autonegotiation must be enabled.

NOTE: On EX Series switches, an interface configuration that disables autonegotiation and manually sets the link speed to 1 Gbps is accepted when you commit the configuration; however, if the interface you are configuring is a Tri-Rate Ethernet copper interface, the configuration is ignored as invalid and autonegotiation is enabled by default.

To correct the invalid configuration and disable autonegotiation:

1. Delete the **no-auto-negotiation** statement and commit the configuration.
2. Set the link speed to 10 or 100 Mbps, set **no-auto-negotiation**, and commit the configuration.

On EX Series switches, if the link speed and duplex mode are also configured, the interfaces use the values configured as the desired values in the negotiation. If autonegotiation is disabled, the link speed and link mode must be configured.

NOTE: On T4000 routers, the **auto-negotiation** command is ignored for interfaces other than Gigabit Ethernet.

Copyright © 2018, Juniper Networks, Inc.

NOTE: On ACX Series routers, when you configure fiber interfaces (fiber media mode) to operate at 1 Gbps, you need to always enable autonegotiation (**auto-negotiation**) to negotiate the speed and duplex settings. You cannot disable autonegotiation (**no-auto-negotiation**) in the fiber media mode. In copper interfaces (copper media mode), autonegotiation is enabled by default. To disable autonegotiation, you need to explicitly configure the link speed to 10 or 100 Mbps, set **no-auto-negotiation**, and commit the configuration.

Default Autonegotiation is automatically enabled. No explicit action is taken after the autonegotiation is complete or if the negotiation fails.

Options **remote-fault (local-interface-online | local-interface-offline)**—(Optional) For M Series, MX Series, T Series, TX Matrix routers, and ACX Series routers only, manually configure remote fault on an interface.

Default: local-interface-online

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

Related Documentation

- *Gigabit Ethernet Autonegotiation Overview*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure) for EX Series Switches with ELS support*

 Copyright © 2018, Juniper Networks, Inc.

backup-destination

Syntax backup-destination *address*;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description For tunnel interfaces, specify the remote address of the backup tunnel.

Options ***address***—Address of the remote side of the connection.

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*
- destination (Tunnels) on page 529

backup-interface

Syntax backup-interface es-*fpc/pic/port*;

Hierarchy Level [edit interfaces es-*fpc/pic/port* es-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure a backup ES Physical Interface Card (PIC). If the primary ES PIC fails, the backup becomes active, inherits all the tunnels and security associations (SAs), and acts as the new next hop for IP Security (IPsec) traffic.

Options **es-*fpc/pic/port***—Name of ES interface to serve as the backup.

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2018, Juniper Networks, Inc.

backup-options

Syntax

```
backup-options {
  interface interface-name;
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure an interface to be used as a backup interface if the primary interface goes down. This is used to support ISDN dial backup operation.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Junos OS Interfaces and Routing Configuration Guide*

Copyright © 2018, Juniper Networks, Inc.

bandwidth (Interfaces)

Syntax bandwidth *rate*;

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

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure the bandwidth value for an interface. This statement is valid for all logical interface types except multilink and aggregated interfaces.

NOTE: We recommend that you be careful when setting this value. Any interface bandwidth value that you configure using the **bandwidth** 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:

cost = reference-bandwidth/bandwidth,

where bandwidth is the physical interface speed. However, if you specify a value for bandwidth using the **bandwidth** statement, that value is used to calculate the interface cost, rather than the actual physical interface bandwidth.

Options *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.

Range: Not limited.

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

Related Documentation

- Configuring the Interface Bandwidth on page 189

Copyright © 2018, Juniper Networks, Inc.

bandwidth-limit (Hierarchical Policer)

Syntax bandwidth-limit *bps*;

Hierarchy Level [edit dynamic-profiles *profile-name* firewall hierarchical-policer aggregate if-exceeding],
[editdynamic-profiles *profile-name* firewall hierarchical-policer premium if-exceeding],
[edit firewall hierarchical-policer aggregate if-exceeding],
[edit firewall hierarchical-policer premium if-exceeding]

Release Information Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles ... if-exceeding]** hierarchy level introduced in Junos OS Release 11.4.

Description On M40e, M120, and M320 (with FFPC and SFPC) edge routers; on MPCs hosted on MX Series routers; on T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs; and on T4000 routers with Type 5 FPC and Enhanced Scaling Type 4 FPC, configure the maximum average bandwidth for premium or aggregate traffic in a hierarchical policer.

Options *bps*—You can specify the number of bits per second either as a decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Range:

- 32,000 through 50,000,000,000 on M Series routers
- 32,000 through 100,000,000,000 on T Series routers
- 32,000 through 18,446,744,073,709,551,615 on MX Series routers

NOTE: When you specify a numeric value beyond the supported bandwidth of the PFE, the router caps the bandwidth at the maximum supported bandwidth of the PFE.

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

Related Documentation

- *Hierarchical Policer Configuration Overview*
- *Policer Bandwidth and Burst-Size Limits*
- *Policer Color-Marking and Actions*
- *Single Token Bucket Algorithm*
- *Determining Proper Burst Size for Traffic Policers*
- *aggregate (Hierarchical Policer)*

 Copyright © 2018, Juniper Networks, Inc.

- burst-size-limit (Hierarchical Policer) on page 469
- premium (Hierarchical Policer) on page 898

bandwidth-limit (Policer for Gigabit Ethernet Interfaces)

Syntax bandwidth-limit *bps*;

Hierarchy Level [edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile policer *cos-policer-name* aggregate],
[edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile policer *cos-policer-name* premium]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define a policer to apply to nonpremium traffic.

Options *bps*—Bandwidth limit, in bits per second. Specify either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Range: 32 Kbps through 32 gigabits per second (Gbps). For IQ2 and IQ2-E interfaces 65,536 bps through 1 Gbps. For 10-Gigabit IQ2 and IQ2-E interfaces 65,536 bps through 10 Gbps.

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

Related Documentation

- *Configuring Gigabit Ethernet Policers*
- burst-size-limit (Policer for Gigabit Ethernet Interfaces) on page 470

Copyright © 2018, Juniper Networks, Inc.

bearer-bandwidth-limit

Syntax bearer-bandwidth-limit *kilobits-per-second*;

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

Release Information Statement introduced in Junos OS Release 8.2.

Description (J4350 and J6350 Services Routers supporting voice over IP with the TGM550 media gateway module) For Fast Ethernet and Gigabit Ethernet interfaces, ISDN BRI interfaces, and serial interfaces with PPP or Frame Relay encapsulation, configure the bearer bandwidth limit (BBL). BBL is used for dynamic call admission control (dynamic CAC) to provide enhanced control over WAN bandwidth.

Options *kilobits-per-second*—The bearer bandwidth limit to be reported to a TGM550 media gateway module, in kilobits per second (kbps).

Range: 0 through 9999 kbps

Default: 1 (dynamic CAC is not enabled 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

- *Junos OS Services Interfaces Library for Routing Devices*
- *Junos OS Interfaces and Routing Configuration Guide*

 Copyright © 2018, Juniper Networks, Inc.

bert-algorithm

Syntax bert-algorithm *algorithm*;

Hierarchy Level [edit interfaces ce1-*fpc/pic/port*],
[edit interfaces ct1-*fpc/pic/port*],
[edit interfaces *interface-name* ds0-options],
[edit interfaces *interface-name* e1-options],
[edit interfaces *interface-name* e3-options],
[edit interfaces *interface-name* t1-options],
[edit interfaces *interface-name* t3-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers.

Description Configure the pattern to send in the bit stream during a bit error rate test (BERT). Applies to T1, E3, T3, and multichannel DS3 interfaces, the channelized interfaces (DS3, OC12, STM1), and channelized IQ and IQE interfaces (E1, E3 and DS3).

NOTE: When configuring CE1 or CT1 interfaces on 10-port Channelized E1/T1 IQE PICs, the **bert-algorithm** statement must be included at the **[edit interfaces ce1-*fpc/pic/port*]** or **[edit interfaces ct1-*fpc/pic/port*]** hierarchy level as appropriate.

Options *algorithm*—Pattern to send in the bit stream. There are two categories of test patterns: pseudorandom and repetitive. Both patterns conform to CCITT/ITU O.151, O.152, O.153, and O.161 standards. The algorithm can be one of the following patterns:

- **all-ones-repeating**—Pattern is all ones.
- **all-zeros-repeating**—Pattern is all zeros.
- **alternating-double-ones-zeros**—Pattern is alternating pairs of ones and zeros.
- **alternating-ones-zeros**—Pattern is alternating ones and zeros.
- **pseudo-2e3**—Pattern is $2^3 – 1$.
- **pseudo-2e4**—Pattern is $2^4 – 1$.
- **pseudo-2e5**—Pattern is $2^5 – 1$.
- **pseudo-2e6**—Pattern is $2^6 – 1$.
- **pseudo-2e7**—Pattern is $2^7 – 1$.
- **pseudo-2e9-o153**—Pattern is $2^9 – 1$, as defined in the O153 standard.
- **pseudo-2e10**—Pattern is $2^{10} – 1$.
- **pseudo-2e11-o152**—Pattern is $2^{11} – 1$, as defined in the O152 standard.

Copyright © 2018, Juniper Networks, Inc.

- **pseudo-2e15-o151**—Pattern is $2^{15} - 1$, as defined in the O151 standard.
- **pseudo-2e17**—Pattern is $2^{17} - 1$.
- **pseudo-2e18**—Pattern is $2^{18} - 1$.
- **pseudo-2e20-o151**—Pattern is $2^{20} - 1$, as defined in the O151 standard.
- **pseudo-2e20-o153**—Pattern is $2^{20} - 1$, as defined in the O153 standard.
- **pseudo-2e21**—Pattern is $2^{21} - 1$.
- **pseudo-2e22**—Pattern is $2^{22} - 1$.
- **pseudo-2e23-o151**—Pattern is $2^{23} - 1$, as defined in the O151 standard.
- **pseudo-2e25**—Pattern is $2^{25} - 1$.
- **pseudo-2e28**—Pattern is $2^{28} - 1$.
- **pseudo-2e29**—Pattern is $2^{29} - 1$.
- **pseudo-2e31**—Pattern is $2^{31} - 1$.
- **pseudo-2e32**—Pattern is $2^{32} - 1$.
- **repeating-1-in-4**—One bit in four is set to 1; the others are set to 0.
- **repeating-1-in-8**—One bit in eight is set to 1; the others are set to 0.
- **repeating-3-in-24**—Three bits in twenty four are set to 1; the others are set to 0.

Default: pseudo-2e3

Required Privilege Level

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

Related Documentation

- *Interface Diagnostics*
- *Configuring E1 BERT Properties*
- *Configuring E3 BERT Properties*
- *Configuring T1 BERT Properties*
- *Configuring T3 BERT Properties*
- *Examples: Configuring T3 Interfaces*
- bert-error-rate on page 461
- bert-period on page 463

Copyright © 2018, Juniper Networks, Inc.

bert-error-rate

Syntax bert-error-rate *rate*;

Hierarchy Level [edit interfaces ce1-*fpc/pic/port*],
[edit interfaces ct1-*fpc/pic/port*],
[edit interfaces *interface-name* ds0-options],
[edit interfaces *interface-name* e1-options],
[edit interfaces *interface-name* e3-options],
[edit interfaces *interface-name* t1-options],
[edit interfaces *interface-name* t3-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers.

Description Configure the bit error rate to use in a BERT procedure. Applies to E1, E3, T1, or T3 interfaces, and to the channelized interfaces (DS3, OC3, OC12, and STM1).

NOTE: When configuring CE1 or CT1 interfaces on 10-port Channelized E1/T1 IQE PICs, the **bert-error-rate** statement must be included at the **[edit interfaces ce1-*fpc/pic/port*]** or **[edit interfaces ct1-*fpc/pic/port*]** hierarchy level as appropriate.

Options *rate*—Bit error rate.

Range: 0 through 7, which corresponds to 10^{-1} (1 error per bit) to 10^{-7} (1 error per 10 million bits)

Default: 0

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

Related Documentation

- bert-algorithm on page 459
- bert-period on page 463
- ds0-options on page 547
- e1-options on page 568
- e3-options on page 569
- t1-options on page 1003
- t3-options on page 1007
- *Interface Diagnostics*
- *Configuring E1 BERT Properties*

Copyright © 2018, Juniper Networks, Inc.

- *Configuring E3 BERT Properties*
- *Configuring T1 BERT Properties*
- *Configuring T3 BERT Properties*
- *Examples: Configuring T3 Interfaces*

Copyright © 2018, Juniper Networks, Inc.

bert-period

Syntax bert-period *seconds*;

Hierarchy Level [edit interfaces ce1-*fpc/pic/port*],
[edit interfaces ct1-*fpc/pic/port*],
[edit interfaces *interface-name* ds0-options],
[edit interfaces *interface-name* e1-options],
[edit interfaces *interface-name* e3-options],
[edit interfaces *interface-name* t1-options],
[edit interfaces *interface-name* t3-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers.

Description Configure the duration of a BERT test. Applies to E1, E3, T1, and T3 interfaces, and to E1, E3, T1, and T3 partitions on the channelized interfaces (CE1, CT1, DS3, OC3, OC12, OC48, STM1, STM4, and STM16).

E1 and T1 IQ, IQE, and standard interfaces support an extended BERT period range, up to 86,400 seconds (24 hours).

NOTE: When configuring CE1 or CT1 interfaces on 10-port Channelized E1/T1 IQE PICs, the **bert-period** statement must be included at the **[edit interfaces ce1-*fpc/pic/port*]** or **[edit interfaces ct1-*fpc/pic/port*]** hierarchy level as appropriate.

Options *seconds*—Test duration. Range and default values vary by interface type.

Range:

- PIC-dependent—Normal BERT period: either 1 through 239 seconds or 1 through 240 seconds
- PIC-dependent—Extended BERT period: from 1 through 86,400 seconds

Default:

- Normal BERT period: 10 seconds
- Extended BERT period (on supported E1 interfaces): 10 seconds
- Extended BERT period (on supported T1 interfaces): 240 seconds

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

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- *Interface Diagnostics*
- *Configuring E1 BERT Properties*
- *Configuring E3 BERT Properties*
- *Configuring T1 BERT Properties*
- *Configuring T3 BERT Properties*
- bert-algorithm on page 459
- bert-error-rate on page 461

bridge-domain

Syntax

```
bridge-domain name;
  vlan-id [ vlan-identifiers ];
}
```

Hierarchy Level

[edit protocols oam ethernet connectivity-fault-management maintenance-domain *maintenance-domain-name*],
[edit protocols oam ethernet connectivity-fault-management maintenance-domain *maintenance-domain-name* virtual-switch *virtual-switch-name*]

Release Information

Statement introduced in Junos OS Release 9.4.

Description

(MX Series routers only) Specify the OAM Ethernet CFM maintenance domain bridge domain.

Options

name—Specify the name of the bridge domain.

vlan-identifiers—Specify one or more VLAN identifiers.

Required Privilege Level

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

Related Documentation

- *Configuring Maintenance Intermediate Points (MIPs)*
- maintenance-domain on page 767

Copyright © 2018, Juniper Networks, Inc.

broadcast

Syntax broadcast *address*;

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

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Set the broadcast address on the network or subnet. On a subnet you cannot specify a host address of 0 (0.0.0.0), nor can you specify a broadcast address (255.255.255.255). For example, in the statement **set interface ge-0/0/0 unit 0 family inet address 10.1.1.0/24**, the subnet address **10.1.1.0** has the host address of 0. Hence, you cannot configure this address. Similarly, for the subnet, you cannot use the broadcast address **10.1.1.255/24**.

Default The default broadcast address has a host portion of all ones.

Options *address*—Broadcast address. The address must have a host portion of either all ones or all zeros. You cannot specify the addresses **0.0.0.0** or **255.255.255.255**.

NOTE: **The edit logical-systems hierarchy is not available on QFabric systems.**

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

Related Documentation

- Configuring the Interface Address on page 207

buildout (E3 or T3 over ATM Interfaces)

Syntax buildout *feet*;

Hierarchy Level [edit interfaces at-*fpc/pic/port* e3-options],
[edit interfaces at-*fpc/pic/port* t3-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For E3 and T3 traffic over ATM interfaces, set the buildout value.

Options *feet*—The buildout value in feet.

Range: 0 through 450 feet (137 meters)

Default: 10 feet (3 meters)

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

Related Documentation

- *Configuring E3 and T3 Parameters on ATM Interfaces*

 Copyright © 2018, Juniper Networks, Inc.

buildout (T1 Interfaces)

Syntax buildout *value*;

Hierarchy Level [edit interfaces ct1-*fpc/pic/port*]
[edit interfaces *interface-name* t1-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers.

Description For T1 interfaces, set the buildout value.

NOTE: When configuring CT1 interfaces on 10-port Channelized E1/T1 IQE PICs and 16-Port Channelized E1/T1 Circuit Emulation MICs, the **buildout** statement must be included at the hierarchy level.

Default The default buildout value is 0 through 132 feet.

Options You can set the buildout value to one of the following:

- **0-132**—0 through 132 feet (0 through 40 meters)
- **133-265**—133 through 265 feet (40 through 81 meters)
- **266-398**—266 through 398 feet (81 through 121 meters)
- **399-531**—399 through 531 feet (121 through 162 meters)
- **532-655**—532 through 655 feet (162 through 200 meters)
- **long-7.5db**—For MX Series only, long buildout with 7.5 dB transmit attenuation
- **long-15db**—For MX Series only, long buildout with 15 dB transmit attenuation
- **long-22.5db**—For MX Series only, long buildout with 22.5 dB transmit attenuation

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

Related Documentation

- *Configuring the T1 Buildout*
- *Junos OS Interfaces and Routing Configuration Guide*

Copyright © 2018, Juniper Networks, Inc.

bundle

Syntax bundle (ml-*fpc/pic/port* | ls-*fpc/pic/port*);

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Associate the multilink interface with the logical interface it is joining. You can include this statement for the **mlfr-end-to-end** and **mlfr-uni-nni** protocol families only.

NOTE:
For M Series routers and T Series routers, the following caveats apply:

- Maximum supported throughput on the bundle interfaces is 45 Mbps.
- Bundling of the logical interfaces under a T3 physical interface into the same or different bundles is not supported.

Options **ml-*fpc/pic/port***—Name of the multilink interface you are linking.

ls-*fpc/pic/port*—Name of the link services interface you are linking.

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

 Copyright © 2018, Juniper Networks, Inc.

burst-size-limit (Hierarchical Policer)

Syntax burst-size-limit *bytes*;

Hierarchy Level [edit dynamic-profiles *profile-name* firewall hierarchical-policer aggregate if-exceeding],
[editdynamic-profiles *profile-name* firewall hierarchical-policer premium if-exceeding],
[edit firewall hierarchical-policer aggregate if-exceeding],
[edit firewall hierarchical-policer premium if-exceeding]

Release Information Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles ... if exceeding]** hierarchy level introduced in Junos OS Release 11.4.

Description On M40e, M120, and M320 (with FFPC and SFPC) edge routers; on MPCs hosted on MX Series routers; on T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs; and on T4000 routers with Type 5 FPC and Enhanced Scaling Type 4 FPC, configure the burst-size limit for premium or aggregate traffic in a hierarchical policer.

Options *bytes*—Burst-size limit in bytes. The minimum recommended value is the maximum transmission unit (MTU) of the IP packets being policed. You can specify the value either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Range: 1500 through 2,147,450,880 (1500 through 100,000,000,000 on MPCs hosted on MX Series routers)

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

Related Documentation

- *Hierarchical Policer Configuration Overview*
- *Policer Bandwidth and Burst-Size Limits*
- *Policer Color-Marking and Actions*
- *Single Token Bucket Algorithm*
- *Determining Proper Burst Size for Traffic Policers*
- *Hierarchical Policers*
- *aggregate (Hierarchical Policer)*
- bandwidth-limit (Hierarchical Policer) on page 456
- premium (Hierarchical Policer) on page 898

burst-size-limit (Policer for Gigabit Ethernet Interfaces)

Syntax burst-size-limit *bytes*;

Hierarchy Level [edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile policer *cos-policer-name* aggregate],
[edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile policer *cos-policer-name* premium]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define a policer to apply to nonpremium traffic.

Options *bytes*—Burst length.

Range: 1500 through 100,000,000 bytes

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

Related Documentation

- *Configuring Gigabit Ethernet Policers*
- bandwidth-limit (Policer for Gigabit Ethernet Interfaces) on page 457

 Copyright © 2018, Juniper Networks, Inc.

byte-encoding

Syntax byte-encoding (nx56 | nx64);

Hierarchy Level [edit interfaces t1-*fpc/pic/port*],
[edit interfaces *interface-name* ds0-options],
[edit interfaces *interface-name* t1-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers.

Description Set the byte encoding on a DS0 or T1 interface to use 7 bits per byte or 8 bits per byte.

NOTE: When configuring T1 interfaces on the 10-port Channelized E1/T1 IQE PIC, the **byte-encoding** statement must be included at the **[edit interfaces t1-*fpc/pic/port*]** hierarchy level.

Default The default byte encoding is 8 bits per byte (nx64).

Options **nx56**—Use 7 bits per byte.

nx64—Use 8 bits per byte.

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

Related Documentation
- *Configuring T1 Byte Encoding*

Copyright © 2018, Juniper Networks, Inc.

bytes

Syntax

```
bytes {
  c2 value;
  e1-quiet value;
  f1 value;
  f2 value;
  s1 value;
  z3 value;
  z4 value;
}
```

Hierarchy Level [edit interfaces *interface-name* sonet-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Set values in some SONET/SDH header bytes.

Options

c2 ***value***—Path signal label SONET/SDH overhead byte. SONET/SDH frames use the C2 byte to indicate the contents of the payload inside the frame. SONET/SDH interfaces use the C2 byte to indicate whether the payload is scrambled.

Range: 0 through 255

Default: **0xCF**

e1-quiet ***value***—Default idle byte sent on the orderwire SONET/SDH overhead bytes. The router does not support the orderwire channel, and hence sends this byte continuously.

Range: 0 through 255

Default: 0x7F

f1 ***value***, **f2** ***value***, **z3** ***value***, **z4** ***value***—SONET/SDH overhead bytes.

Range: 0 through 255

Default: 0x00

s1 ***value***—Synchronization message SONET overhead byte. This byte is normally controlled as a side effect of the system reference clock configuration and the state of the external clock coming from an interface if the system reference clocks have been configured to use an external reference.

Range: 0 through 255

Default: 0xCC

Required Privilege Level

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

Related Documentation

- *Configuring SONET/SDH Header Byte Values to Identify Error Conditions*
- *no-concatenate*

 Copyright © 2018, Juniper Networks, Inc.

calculation-weight

Syntax

```
calculation-weight {
  delay delay-value;
  delay-variation delay-variation-value;
}
```

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles *profile-name*]

Release Information Statement introduced in Junos OS Release 11.1.
Statement introduced in Junos OS Release 11.4 for EX Series switches.

Description Configure the calculation weight for delay and delay variation.

NOTE: This option is applicable only for two-way delay measurement.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- *Configuring an Iterator Profile*
- *Configuring an Iterator Profile on a Switch (CLI Procedure)*
- *delay*
- *delay-variation*

callback

Syntax callback;

Hierarchy Level [edit interfaces dl*n* unit *logical-unit-number* dialer-options incoming-map],
[edit logical-systems *logical-system-name* interfaces dl*n* unit *logical-unit-number* dialer-options incoming-map]

Release Information Statement introduced in Junos OS Release 7.5.

Description On J Series Services Routers with interfaces configured for ISDN, configure the dialer to terminate the incoming call and call back the originator after the callback wait period. The default wait time is 5 seconds. To configure the wait time, include the **callback-wait-period** statement at the **[edit interfaces dl *n* unit *logical-unit-number* dialer-options]** hierarchy level.

NOTE: The **incoming-map** statement is mandatory for the router to accept any incoming ISDN calls.

If the **callback** statement is configured, you cannot use the **caller *caller-id*** statement at the **[edit interfaces dl*n* unit *logical-unit-number* dialer-options]** hierarchy level.

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

Related Documentation

- *Junos OS Interfaces and Routing Configuration Guide*
- callback-wait-period on page 475

 Copyright © 2018, Juniper Networks, Inc.

callback-wait-period

Syntax callback-wait-period *time*;

Hierarchy Level [edit interfaces dl*n* unit *logical-unit-number* dialer-options],
[edit logical-systems *logical-system-name* interfaces dl*n* unit *logical-unit-number* dialer-options]

Release Information Statement introduced in Junos OS Release 7.5.

Description On J Series Services Routers with interfaces configured for ISDN with callback, specify the amount of time the dialer waits before calling back the caller. The default wait time is 5 seconds. The wait time is necessary because, when a call is rejected, the switch waits for up to 4 seconds on point-to-multipoint connections to ensure no other device accepts the call before sending the DISCONNECT message to the originator of the call. However, the default time of 5 seconds may not be sufficient for different switches or may not be needed on point-to-point connections.

To configure callback mode, include the **callback** statement at the **[edit interfaces dl*n* unit *logical-unit-number* dialer-options]** hierarchy level.

If the **callback** statement is configured, you cannot use the **caller *caller-id*** statement at the **[edit interfaces dl*n* unit *logical-unit-number* dialer-options]** hierarchy level.

Options *time*—Time the dialer waits before calling back the caller.

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

Related Documentation
- *Junos OS Interfaces and Routing Configuration Guide*

Copyright © 2018, Juniper Networks, Inc.

caller

Syntax caller (*caller-id* | accept-all);

Hierarchy Level [edit interfaces dl*n* unit *logical-unit-number* dialer-options incoming-map],
[edit logical-systems *logical-system-name* interfaces dl*n* unit *logical-unit-number* dialer-options incoming-map]

Release Information Statement introduced in Junos OS Release 7.5.

Description On J Series Services Routers with interfaces configured for ISDN, specify the dialer to accept a specified caller number or accept all incoming calls.

Options ***caller-id***—Incoming caller number. You can configure multiple caller IDs on a dialer. The caller ID of the incoming call is matched against all caller IDs configured on all dialers. The dialer matching the caller ID is looked at for further processing. Only a precise match is a valid match, For example, the configured caller ID 1-222-333-4444 or 222-333-4444 will match the incoming caller ID 1-222-333-4444.

If the incoming caller ID has fewer digits than the number configured, it is not a valid match. Duplicate caller IDs are not allowed on different dialers; however, for example, the numbers 1-408-532-1091, 408-532-1091, and 532-1091 can still be configured on different dialers.

Only one B-channel can map to one dialer. If one dialer is already mapped, any other call mapping to the same dialer is rejected (except in the case of a multilink dialer). If no dialer caller is configured on a dialer, that dialer will not accept any calls.

accept-all—Any incoming call in an associated interface is accepted.

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

Related Documentation

- *Junos OS Interfaces and Routing Configuration Guide*

 Copyright © 2018, Juniper Networks, Inc.

calling-number

Syntax calling-number *number*;

Hierarchy Level [edit interfaces br-*pim*/0/*port* isdn-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description On J Series Services Routers with ISDN interfaces, configure the calling number to include in outgoing calls.

Options ***number***—Calling number.

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

Related Documentation

- *Configuring ISDN Physical Interface Properties*
- *Junos OS Interfaces and Routing Configuration Guide*

cbit-parity

Syntax (cbit-parity | no-cbit-parity);

Hierarchy Level [edit interfaces *interface-name* t3-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For T3 interfaces only, enable or disable C-bit parity mode, which controls the type of framing that is present on the transmitted T3 signal. When C-bit parity mode is enabled, the C-bit positions are used for the far-end block error (FEBE), far-end alarm and control (FEAC), terminal data link, path parity, and mode indicator bits, as defined in ANSI T1.107a-1989. For ATM and ATM2 IQ2 and IQ2-E interfaces, M23 framing is used when the **no-cbit-parity** statement is included. For all other interfaces, M13 framing is used when the **no-cbit-parity** statement is included.

Default C-bit parity mode is enabled.

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

Related Documentation

- *Configuring E3 and T3 Parameters on ATM Interfaces*
- *Disabling T3 C-Bit Parity Mode*

Copyright © 2018, Juniper Networks, Inc.

cbr

Syntax cbr *rate*;

Hierarchy Level [edit interfaces at-*fpc/pic/port* atm-options vpi *vpi-identifier* shaping],
[edit interfaces at-*fpc/pic/port* unit *logical-unit-number* address *address* family *family* multipoint-destination *address* shaping],
[edit interfaces at-*fpc/pic/port* unit *logical-unit-number* shaping],
[edit logical-systems *logical-system-name* interfaces at-*fpc/pic/port* unit *logical-unit-number* address *address* family *family* multipoint-destination *address* shaping],
[edit logical-systems *logical-system-name* interfaces at-*fpc/pic/port* unit *logical-unit-number* shaping]

Release Information Statement introduced before Junos OS Release 7.4.

Description For ATM encapsulation only, define a constant bit rate bandwidth utilization in the traffic-shaping profile.

Default Unspecified bit rate (UBR); that is, bandwidth utilization is unlimited.

Options ***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.

For ATM1 and ATM2 OC3 interfaces, the maximum available rate is 100 percent of ***line-rate***, or 135,600,000 bps. For ATM1 OC12 interfaces, the maximum available rate is 50 percent of ***line-rate***, or 271,263,396 bps. For ATM2 IQ interfaces, the maximum available rate is 542,526,792 bps.

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

Related Documentation

- *Defining the ATM Traffic-Shaping Profile Overview*
- rtvbr on page 944
- shaping on page 965
- vbr on page 1071

 Copyright © 2018, Juniper Networks, Inc.

cell-bundle-size

Syntax cell-bundle-size *cells*;

Hierarchy Level [edit interfaces at-*fpc/pic/port* atm-options],
[edit interfaces at-*fpc/pic/port* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces at-*fpc/pic/port* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For ATM2 IQ interfaces using ATM Layer 2 circuit cell-relay transport mode only, configure the maximum number of ATM cells per frame.

NOTE: For MIC-3D-8OC3-2OC12-ATM on MX104 routers, ensure that the configured **cell-bundle-size** is less than 30 for an ATM interface that is configured with **atm-ccc-cell-relay** encapsulation. If the configured **cell-bundle-size** is greater than or equal to 30 and the traffic is passing through the interface at line rate, it might lead to AFEB crash.

Options *cells*—Maximum number of cells.

Default: 1 cell

Range: 1 through 176 cells

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

Related Documentation

- *Configuring the Layer 2 Circuit Cell-Relay Cell Maximum Overview*

Copyright © 2018, Juniper Networks, Inc.

chap

Syntax

```
chap {
  access-profile name;
  challenge-length minimum minimum-length maximum maximum-length;
  default-chap-secret name;
  local-name name;
  passive;
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Allow each side of a link to challenge its peer, using a “secret” known only to the authenticator and that peer. The secret is not sent over the link.

By default, PPP CHAP is disabled. If CHAP is not explicitly enabled, the interface makes no CHAP challenges and denies all incoming CHAP challenges.

For ATM2 IQ interfaces only, you can configure CHAP on the logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:

- **atm-ppp-llc**—PPP over AAL5 LLC encapsulation.
- **atm-ppp-vc-mux**—PPP over AAL5 multiplex encapsulation.

BEST PRACTICE: **On inline service (si) interfaces for L2TP, only the chap statement itself is typically used for subscriber management. We recommend that you leave the subordinate statements at their default values.**

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

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

Related Documentation

- Configuring the PPP Challenge Handshake Authentication Protocol on page 136
- *Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile*
- *Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface*

Copyright © 2018, Juniper Networks, Inc.

chap-secret

Syntax chap-secret *chap-secret*;

Hierarchy Level [edit access profile *profile-name* client *client-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For interfaces with PPP encapsulation on which the PPP Challenge Handshake Authentication Protocol (CHAP) is configured, configure the shared secret (the CHAP secret key associated with a peer), as defined in RFC 1994.

NOTE: This statement is not supported for L2TP LNS on MX Series routers.

Options *chap-secret*—The secret key associated with a peer.

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

Related Documentation

- *Configuring the CHAP Secret for an L2TP Profile*
- Configuring PPP CHAP Authentication on page 196
- **pap-password on page 853**
- *Junos OS Administration Library*

circuit-type

Syntax circuit-type;

Hierarchy Level [edit interfaces *interface-name* auto-configure vlan-ranges authentication username-include],
[edit interfaces *interface-name* auto-configure stacked-vlan-ranges authentication username-include],

Release Information Statement introduced in Junos OS Release 10.0.

Description Specify that the circuit type is concatenated with the username during the subscriber authentication process.

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

Related Documentation
- *Configuring VLAN Interface Username Information for AAA Authentication*

cisco-interoperability

Syntax cisco-interoperability send-lip-remove-link-for-link-reject;

Hierarchy Level [edit interfaces *interface-name* mlfr-uni-nni-bundle-options]

Release Information Statement introduced in Junos OS Release 7.4.

Description FRF.16 interoperability settings.

Options **send-lip-remove-link-for-link-reject**—Send Link Integrity Protocol remove link when an add-link rejection message is received.

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

Related Documentation
- *Junos OS Services Interfaces Library for Routing Devices*

 Copyright © 2018, Juniper Networks, Inc.

classifier

Syntax

```
classifier {
  per-unit-scheduler {
    forwarding-class class-name {
      loss-priority (high | low);
    }
  }
}
```

Hierarchy Level [edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile output-priority-map]

Release Information Statement introduced before Junos OS Release 7.4.

Description For Gigabit Ethernet IQ and 10-Gigabit Ethernet interfaces only, define the classifier for the output priority map to be applied to outgoing frames on this interface.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *Specifying an Output Priority Map*
- input-priority-map on page 677

clear-dont-fragment-bit

Syntax clear-dont-fragment-bit;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Clear the don't-fragment (DF) bit on all IP version 4 (IPv4) packets entering a generic routing encapsulation (GRE) tunnel. If the encapsulated packet's size exceeds the tunnel's maximum transmission unit (MTU), the packet is fragmented before encapsulation. The statement is supported only on MX Series routers and all M Series routers except the M320 router.

When you configure the **clear-dont-fragment-bit** statement on an interface with the MPLS protocol family enabled, you must specify an MTU value. This MTU value must not be greater than maximum supported value, which is 9192.

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

 Copyright © 2018, Juniper Networks, Inc.

clock-rate

Syntax clock-rate *rate*;

Hierarchy Level [edit interfaces *interface-name* serial-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For EIA-530 and V.35 interfaces, configure the interface speed, in megahertz (MHz).

Options *rate*—You can specify one of the following rates:

- 2.048 MHz
- 2.341 MHz
- 2.731 MHz
- 3.277 MHz
- 4.096 MHz
- 5.461 MHz
- 8.192 MHz
- 16.384 MHz

Default: 16.384 MHz

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

Related Documentation

- Configuring the Serial Clocking Mode on page 342

clocking

Syntax clocking (external [interface *interface-name*] | internal);

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
interface option added in Junos OS Release 8.2.
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers.

Description For interfaces that can use various clock sources, configure the source of the transmit clock on each interface.

NOTE: On Channelized SONET/SDH PICs, if you set the parent (or the master) controller clock to **external**, then you must set the child controller clocks to the default value—that is, **internal**.

For example, on the Channelized STM1 PIC, if the clock on the Channelized STM1 interface (which is the master controller) is set to **external**, then you must not configure the CE1 interface (which is the child controller) clock to **external**. Instead you must configure the CE1 interface clock to **internal**.

Options **external**—The clock source is provided by the data communication equipment (DCE).

interface ***interface-name***—Configure clocking for the drop-and insert feature. When configuring this feature, both ports must use the same clock source: either the router's internal clock or an external clock on one of the interfaces. If an external clock source is required, one interface must specify clocking external and the other must specify the same clock.

internal—Use the internal stratum 3 clock as the reference clock.

Default: internal

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

Related Documentation

- Configuring the Clock Source on page 128
- *Configuring the Clock Source on SONET/SDH Interfaces*
- *Clock Sources on Channelized Interfaces*
- *Configuring a Channelized T1/E1 Interface to Drop and Insert Time Slots*
- loop-timing on page 750

 Copyright © 2018, Juniper Networks, Inc.

clocking-mode

Syntax clocking-mode (dce | internal | loop);

Hierarchy Level [edit interfaces *interface-name* serial-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For EIA-530 and V.35 interfaces, configure the clock mode. You cannot configure **clocking-mode dce** on a DTE router using an X.21 serial line protocol (detected automatically when an X.21 cable is plugged into the serial interface).

Options **dce**—DCE timing (DTE mode only, not valid for X.21).

internal—Internal baud timing.

loop—Loop timing.

Default: loop

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

Related Documentation

- Configuring the Serial Clocking Mode on page 342

Copyright © 2018, Juniper Networks, Inc.

community (Policy Options)

Syntax

```
community name {
  invert-match;
  members [ community-ids ];
}
```

Hierarchy Level

[edit dynamic policy-options],
[edit logical-systems *logical-system-name* policy-options],
[edit policy-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for configuration in the dynamic database introduced in Junos OS Release 9.5.
Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.
Support for BGP large community introduced in Junos OS Release 17.3 for MX Series, PTX Series, and QFX Series.

Description

Define a community, extended community or large community for use in a routing policy match condition.

Options

name—Name that identifies the regular expression. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters. To include spaces in the name, enclose it in quotation marks (“ ”).

invert-match—Invert the results of the community expression matching. The **community** match condition defines a regular expression and if it matches the community attribute of the received prefix, Junos OS returns a TRUE result. If not, Junos OS returns a FALSE result. The **invert-match** statement makes Junos OS behave to the contrary. If there is a match, Junos OS returns a FALSE result. If there is no match, Junos OS returns a TRUE result.

members *community-ids*—One or more community members. If you specify more than one member, you must enclose all members in brackets.

The format for ***community-ids*** is:

as-number:*community-value*

Starting in Junos OS Release 15.1, you can apply a wildcard member **segmented-nh:.*:0** to apply the BGP policy to all the S-PMSI A-D routes carrying extended community information.

as-number is the AS number and can be a value in the range from 0 through 65,535. ***community-value*** is the community identifier and can be a number in the range from 0 through 65,535.

You also can specify ***community-ids*** for communities as one of the following well-known community names, which are defined in RFC 1997, *BGP Communities Attribute*:

Copyright © 2018, Juniper Networks, Inc.

- **no-export**—Routes containing this community name are not advertised outside a BGP confederation boundary.
- **no-advertise**—Routes containing this community name are not advertised to other BGP peers.
- **no-export-subconfed**—Routes containing this community name are not advertised to external BGP peers, including peers in other members' ASs inside a BGP confederation.

You can explicitly exclude BGP community information with a static route using the **none** option. Include **none** when configuring an individual route in the **route** portion of the **static** statement to override a **community** option specified in the **defaults** portion of the statement.

The format for extended ***community-ids*** is the following:

type:administrator:assigned-number

type is the type of extended community and can be either a **bandwidth**, **target**, **origin**, **domain-id**, **src-as**, or **rt-import** community or a 16-bit number that identifies a specific BGP extended community. The **target** community identifies the destination to which the route is going. The **origin** community identifies where the route originated. The **domain-id** community identifies the OSPF domain from which the route originated. The **src-as** community identifies the autonomous system from which the route originated. The **rt-import** community identifies the route to install in the routing table.

NOTE: **For src-as, you can specify only an AS number and not an IP address. For rt-import, you can specify only an IP address and not an AS number.**

administrator is the administrator. It is either an AS number or an IPv4 address prefix, depending on the type of extended community.

assigned-number identifies the local provider.

The format for linking a bandwidth with an AS number is:

bandwidth:*as-number*:*bandwidth*

as-number specifies the AS number and ***bandwidth*** specifies the bandwidth in bytes per second.

Copyright © 2018, Juniper Networks, Inc.

NOTE: In Junos OS Release 9.1 and later, you can specify 4-byte AS numbers as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*, as well as the 2-byte AS numbers that are supported in earlier releases of the Junos OS. In plain-number format, you can configure a value in the range from 1 through 4,294,967,295. To configure a **target** or **origin** extended community that includes a 4-byte AS number in the plain-number format, append the letter "L" to the end of number. For example, a target community with the 4-byte AS number 334,324 and an assigned number of 132 is represented as **target:334324L:132**.

In Junos OS Release 9.2 and later, you can also use AS-dot notation when defining a 4-byte AS number for the **target** and **origin** extended communities. Specify two integers joined by a period: *16-bit high-order value in decimal*.*16-bit low-order value in decimal*. For example, the 4-byte AS number represented in plain-number format as 65546 is represented in AS-dot notation as 1.10.

As defined in RFC 8092, BGP large community uses 12-byte encoding and the format for BGP large ***community-ids*** is:

large: *global-administrator*:*assigned-number*:*assigned-number*

large indicates BGP large community.

global-administrator is the administrator. It is a 4-byte AS number.

assigned-number is a 4-byte value used to identify the local provider. BGP large community uses two 4-byte assigned number to identify the local provider.

Required Privilege Level

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

Related Documentation

- *Understanding BGP Communities, Extended Communities, and Large Communities as Routing Policy Match Conditions*
- *Understanding How to Define BGP Communities and Extended Communities*
- *dynamic-db*

 Copyright © 2018, Juniper Networks, Inc.

compatibility-mode

Syntax compatibility-mode (adtran | digital-link | kentrox | larscom | verilink) <subrate *value*>;

Hierarchy Level [edit interfaces *interface-name* e3-options],
[edit interfaces *interface-name* t3-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the E3 or T3 interface to be compatible with the channel service unit (CSU) at the remote end of the line.

NOTE: The compatibility-mode statement at the [edit interfaces *interface-name* e3-options] hierarchy level is not valid for IQE PICs.

Default If you omit this option, the full E3 or T3 rate is used.

Options **adtran**—For T3 IQ interfaces only, configure compatibility with Adtran CSUs.

digital-link—Configure compatibility with Digital Link CSUs. If you include this option on an E3 interface, you must also disable payload scrambling.

kentrox—Configure compatibility with Kentrox CSUs. Kentrox subrate is valid for E3 IQ and T3 IQ interfaces only.

larscom—For T3 and T3 IQ interfaces only, configure compatibility with Larscom CSUs.

verilink—For T3 IQ and T3 IQE interfaces only, configure compatibility with Verilink CSUs.

NOTE: Verilink configuration is not functional if an IQ interface is paired with an IQE interface.

subrate *value*—Subrate of the E3 or T3 line.

Range: For Kentrox CSUs on E3 IQ interfaces and T3 IQ interfaces the subrate value must match the value configured on the CSU. Each increment of the subrate value corresponds to a rate increment of about 0.5 Mbps.

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

Related Documentation
- *Configuring the E3 CSU Compatibility Mode*

- *Configuring the T3 CSU Compatibility Mode*
- payload-scrambler on page 823

compression (PPP Properties)

Syntax

```
compression {
  acfc;
  pfc;
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

For interfaces with PPP encapsulation, set Link Control Protocol (LCP) compression options.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- Configuring PPP Address and Control Field Compression on page 149
- Configuring the PPP Protocol Field Compression on page 151

Copyright © 2018, Juniper Networks, Inc.

compression (Voice Services)

Syntax

```
compression {
  rtp {
    f-max-period number;
    queues [ queue-numbers ];
    port {
      minimum port-number;
      maximum port-number;
    }
  }
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure the compression properties for voice services traffic.

The remaining statements are described separately.

Required Privilege Level

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

compression-device

Syntax

compression-device *interface-name*;

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the compression interface for voice services traffic.

Options

interface-name—Logical interface used for compression.

Required Privilege Level

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*
- *Junos OS Interfaces and Routing Configuration Guide*

Copyright © 2018, Juniper Networks, Inc.

connections

Syntax

```
connections {
  interface-switch connection-name {
    interface interface-name.unit-number;
    interface interface-name.unit-number;
  }
}
```

Hierarchy Level [edit protocols]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the connection between two circuits in a circuit cross-connect (CCC) connection.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- Defining the Connection for Switching Cross-Connects on page 281
- *MPLS Applications Feature Guide*

connection-protection-tlv

Syntax connection-protection-tlv;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* continuity-check]

Description Includes connection protection OUI TLV in continuity check messages (CCM) .The TLV is responsible for carrying the flag information within CCM PDUs. Though this OUI TLV will be included in the CCM frames by provider edge devices, the value is updated by the provider routers in case the traffic to the other end of the network is forwarded by the facility protection tunnel .

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

Related Documentation

- **connectivity-fault-management on page 495**
- *Configuring MAC Flush Message Processing in CET Mode*
- *Example: Configuring an Action Profile Based on Connection Protection TLVs*

 Copyright © 2018, Juniper Networks, Inc.

connectivity-fault-management

Syntax

```
connectivity-fault-management {
  action-profile profile-name {
    action {
      interface-down;
      log-and-generate-ais {
        interval(1m | 1s);
        levelvalue;
        priority value;
      }
    }
    default-actions {
      interface-down;
    }
    event {
      ais-trigger-condition {
        adjacency-loss;
        all-defects;
        cross-connect-ccm;
        erroneous-ccm;
        receive-ais;
      }
      adjacency-loss;
      interface-status-tlv (down | lower-layer-down);
      port-status-tlv blocked;
      rdi;
    }
  }
  linktrace {
    age (30m | 10m | 1m | 30s | 10s);
    path-database-size path-database-size;
  }
  maintenance-domain domain-name {
    bridge-domain <vlan-id [ vlan-ids ]>;
    instance routing-instance-name;
    interface interface-name;
    level number;
    name-format (character-string | none | dns | mac+2oct);
    maintenance-association ma-name {
      protect-maintenance-association protect-ma-name;
      remote-maintenance-association remote-ma-name;
      short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
      continuity-check {
        convey-loss-threshold;
        hold-interval minutes;
        interface-status-tlv;
        interval (10m | 10s | 1m | 1s| 100ms);
        loss-threshold number;
        port-status-tlv;
      }
      mep mep-id {
        auto-discovery;
        direction (up | down);
```

```
                interface interface-name (protect | working);
                lowest-priority-defect (all-defects | err-xcon | mac-rem-err-xcon | no-defect |
                  rem-err-xcon | xcon );
                priority number;
                remote-mep mep-id {
                  action-profile profile-name;
                  sla-iterator-profile profile-name {
                    data-tlv-size size;
                    iteration-count count-value;
                    priority priority-value;
                    detect-loc;
                  }
                }
              }
            }
            virtual-switch routing-instance-name {
              bridge-domain name <vlan-ids [ vlan-ids ]>;
            }
          }
          no-aggregate-delegate-processing;
          performance-monitoring {
            delegate-server-processing;
            hardware-assisted-timestamping;
            hardware-assisted-keepalives;
            sla-iterator-profiles {
              profile-name {
                avg-fd-twoway-threshold;
                avg-ifdv-twoway-threshold;
                avg-flr-forward-threshold;
                avg-flr-backward-threshold;
                disable;
                calculation-weight {
                  delay delay-weight;
                  delay-variation delay-variation-weight;
                }
                cycle-time milliseconds;
                iteration-period connections;
                measurement-type (loss | statistical-frame-loss | two-way-delay);
              }
            }
          }
        }
```

Hierarchy Level [edit protocols oam ethernet]

Release Information Statement introduced in Junos OS Release 8.4.

Description For Ethernet interfaces on M7i and M10i routers with Enhanced CFEB (CFEB-E), and on M120, M320, MX Series, and T Series routers, specify connectivity fault management for IEEE 802.1ag Operation, Administration, and Management (OAM) support.
In Junos OS Release 9.3 and later, this statement is also supported on aggregated Ethernet interfaces.

The remaining statements are explained separately. See CLI Explorer.

 Copyright © 2018, Juniper Networks, Inc.

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

Related Documentation
- *IEEE 802.1ag OAM Connectivity Fault Management Overview*

container-devices

Syntax

```
container-devices {
  device-count number;
}
```

Hierarchy Level [edit chassis]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify the container devices configuration. The **number** option specifies the number of sequentially numbered container interfaces, from **ci0** to **ci127** maximum.

Options **number**—Number of container devices.

Range: 1 through 128

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

Related Documentation
- *Displaying APS Using a Container Interface with ATM Encapsulation*
- *Configuring Container Interfaces for APS on SONET Links*

container-list

Syntax container-list [*container-interface-names*];

Hierarchy Level [edit interfaces container-options]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify a list of container interfaces; for example: **ci*0***, **ci*1***, and up to **ci*127***.

Options ***container-interface-names***—Name of each container interface.

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

Related Documentation

- *Displaying APS Using a Container Interface with ATM Encapsulation*
- *Configuring Container Interfaces for APS on SONET Links*
- container-options on page 499

 Copyright © 2018, Juniper Networks, Inc.

container-options

Syntax

```
container-options {
  container-list [ container-interface-names ];
  container-type aps;
  member-interface-type sonet {
    member-interface-speed [ speed ];
  }
}
```

Hierarchy Level [edit interfaces]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify the container interface options.

Options **interface-name**—Name of the SONET or the container interface.

aps—Specify the member link interface type of the container as APS.

sonet—Protocol type of the container interface.

speed—Set interface speed to OC3, OC12, OC48, OC192, OC768, or mixed.

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

Related Documentation

- *Displaying APS Using a Container Interface with ATM Encapsulation*
- *Configuring Container Interfaces for APS on SONET Links*

container-type

Syntax container-type *aps*;

Hierarchy Level [edit interfaces container-options]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify the container-options interface type.

Options **aps**—Configure the interface type to be Automatic Protection Switching (APS).

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

Related Documentation

- *Displaying APS Using a Container Interface with ATM Encapsulation*
- *Configuring Container Interfaces for APS on SONET Links*

 Copyright © 2018, Juniper Networks, Inc.

continuity-check

Syntax

```
continuity-check {
  convey-loss-threshold;
  hold-interval minutes;
  interface-status-tlv;
  interval (10m | 10s | 1m | 1s| 100ms | 10ms);
  loss-threshold number;
  port-status-tlv;
}
```

Hierarchy Level

[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name*]

Release Information

Statement introduced in Junos OS Release 8.4.

Description

Specify continuity check protocol options.

Options

convey-loss-threshold—Enable loss-threshold-tlv transmission.

hold-interval *minutes*—Specify the continuity check hold-interval, in minutes.

interface-status-tlv—Enable interface-status-tlv transmission.

interval *(10m | 10s | 1m | 1s| 100ms | 10ms)*—Specify the continuity check interval.

loss-threshold *minutes*—Specify the loss-threshold, in minutes.

port-status-tlv—Enable port-status-tlv transmission.

Required Privilege Level

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

Related Documentation

- *Configuring Continuity Check Protocol Parameters for Fault Detection*

Copyright © 2018, Juniper Networks, Inc.

control-channel

Syntax

```
control-channel channel-name {
  vlan vlan-id;
  interface name interface-name
}
```

Hierarchy Level [edit protocols protection-group ethernet-ring *name* (east-interface | west-interface)]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

Description Configure the Ethernet RPS control channel logical interface to carry the RAPS PDU. The related physical interface is the physical ring port.

Options **vlan *vlan-id***—If the control channel logical interface is a trunk port, then a dedicated **vlan *vlan-id*** defines the dedicated VLAN channel to carry the RAPS traffic. Only configure the ***vlan-id*** when the control channel logical interface is the trunk port.

interface name *interface-name*—Interface name of the control channel.

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

Related Documentation

- *Ethernet Ring Protection Switching Overview*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

 Copyright © 2018, Juniper Networks, Inc.

control-polarity

Syntax control-polarity (negative | positive);

Hierarchy Level [edit interfaces *interface-name* serial-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For X.21 interfaces only, configure the control signal polarity.

Options **positive**—Positive signal polarity.

negative—Negative signal polarity.

Default: **positive**

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

Related Documentation

- Configuring Serial Signal Polarities on page 347

control-signal

Syntax control-signal (assert | de-assert | normal);

Hierarchy Level [edit interfaces *interface-name* serial-options dce-options],
[edit interfaces *interface-name* serial-options dte-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For X.21 interfaces only, configure the to-DCE signal.

Options **assert**—The to-DCE signal must be asserted.

de-assert—The to-DCE signal must be deasserted.

normal—Normal request-to-send (RTS) signal handling, as defined by ITU-T Recommendation X.21.

Default: normal

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

Related Documentation

- Configuring the Serial Signal Handling on page 344

Copyright © 2018, Juniper Networks, Inc.

copy-tos-to-outer-ip-header

Syntax copy-tos-to-outer-ip-header;

Hierarchy Level [edit interfaces at-*fpc/pic/port* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces at-*fpc/pic/port* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 8.2.

Description For GRE tunnel interfaces only, enable the inner IP header's TOS bits to be copied to the outer IP packet header.

Default If you omit this statement, the TOS bits in the outer IP header are set to 0.

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

Related Documentation

- *Class of Service Feature Guide for Routing Devices and EX9200 Switches*

core-dump

Syntax (core-dump | no-core-dump);

Hierarchy Level [edit interfaces mo-*fpc/pic/port* multiservice-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For monitoring services interfaces only, a useful tool for isolating the cause of a problem. Core dumping is enabled by default. The directory **/var/tmp** contains core files. The Junos OS saves the current core file (0) and the four previous core files, which are numbered 1 through 4 (from newest to oldest):

- **core-dump**—Enable the core dumping operation.
- **no-core-dump**—Disable the core dumping operation.

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

Related Documentation

- Configuring Multiservice Physical Interface Properties on page 170
- *Junos OS Services Interfaces Library for Routing Devices*

 Copyright © 2018, Juniper Networks, Inc.

crc-major-alarm-threshold

Syntax crc-major-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5);

Hierarchy Level [edit interfaces *interface-name* t1-options]

Release Information Statement introduced in Junos OS Release 8.5.
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers.

Description Major alarm error thresholds for T1 CRC errors. When the threshold is exceeded for one second, a defect condition is declared. If the defect condition continues for the monitoring period, an alarm condition is declared.

Default 10-second monitoring period for all settings except **1e-5**. The **1e-5** value uses a 50-second monitoring period.

Options

1e-3—Error rate expressed as the number of errors per number of bits. The value **1e-3** is one crc error in 10^3 bits.

1e-4—Error rate expressed as the number of errors per number of bits. The value **1e-4** is one crc error in 10^4 bits.

1e-5—Error rate expressed as the number of errors per number of bits. The value **1e-5** is one crc error in 10^5 bits.

5e-4—Error rate expressed as the number of errors per number of bits. The value **5e-4** is five crc errors in 10^4 bits.

5e-5—Error rate expressed as the number of errors per number of bits. The value **5e-5** is five crc errors in 10^5 bits.

Default: 5e-5

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

Related Documentation

- *Configuring T1 CRC Error Major Alarm Thresholds*

Copyright © 2018, Juniper Networks, Inc.

crc-minor-alarm-threshold

Syntax crc-minor-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5 | 5e-6 | 1e-6);

Hierarchy Level [edit interfaces *interface-name* t1-options]

Release Information Statement introduced in Junos OS Release 8.5.
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers.

Description Minor alarm error thresholds for T1 CRC errors. When the threshold is exceeded for one second, a defect condition is declared. If the defect condition continues for the monitoring period, an alarm condition is declared.

Default 10-second monitoring period for values **1e-3**, **5e-4**, **1e-4**, and **5e-5**. The **1e-5** value uses a 50-second monitoring period. The **5e-6** value uses a 100-second monitoring period. The **1e-6** value uses a 500-second monitoring period.

Options **1e-3**—Error rate expressed as the number of errors per number of bits. The value **1e-3** is one crc error in 10^{3} bits.

1e-4—Error rate expressed as the number of errors per number of bits. The value **1e-4** is one crc error in 10^{4} bits.

1e-5—Error rate expressed as the number of errors per number of bits. The value **1e-5** is one crc error in 10^{5} bits.

1e-6—Error rate expressed as the number of errors per number of bits. The value **1e-5** is one crc error in 10^{6} bits.

5e-4—Error rate expressed as the number of errors per number of bits. The value **5e-4** is five crc errors in 10^{4} bits.

5e-5—Error rate expressed as the number of errors per number of bits. The value **5e-5** is five crc errors in 10^{5} bits.

5e-6—Error rate expressed as the number of errors per number of bits. The value **5e-5** is five crc errors in 10^{6} bits.

Default: 5e-6

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

Related Documentation
- *Configuring T1 CRC Error Minor Alarm Thresholds*

 Copyright © 2018, Juniper Networks, Inc.

cts

Syntax cts (ignore | normal | require);

Hierarchy Level [edit interfaces *interface-name* serial-options dce-options],
[edit interfaces *interface-name* serial-options dte-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For EIA-530 and V.35 interfaces only, configure the from-DCE signal, clear-to-send (CTS).

Options **ignore**—The from-DCE signal is ignored.

normal—Normal CTS signal handling as defined by the TIA/EIA Standard 530.

require—The from-DCE signal must be asserted.

Default: normal

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

Related Documentation

- Configuring the Serial Signal Handling on page 344

cts-polarity

Syntax cts-polarity (negative | positive);

Hierarchy Level [edit interfaces *interface-name* serial-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure CTS signal polarity.

Options **positive**—Positive signal polarity.

negative—Negative signal polarity.

Default: positive

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

Related Documentation

- Configuring Serial Signal Polarities on page 347

Copyright © 2018, Juniper Networks, Inc.

cycle-time

Syntax cycle-time *cycle-time-value*;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles *profile-name*]

Release Information Statement introduced in Junos OS Release 11.1.
Statement introduced in Junos OS Release 11.4 for EX Series switches.

Description Configure the time (in milliseconds) taken between back-to-back transmissions of SLA frames for a single connection.

Options ***cycle-time-value***—Cycle time value in milliseconds.

Range: 10 through 3,600,000

Default: 1000

Required Privilege Level Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- *Configuring an Iterator Profile*
- *Configuring an Iterator Profile on a Switch (CLI Procedure)*

 Copyright © 2018, Juniper Networks, Inc.

damping (Interfaces)

Syntax

```
damping {
  enable;
  half-life seconds;
  max-suppress seconds;
  reuse number;
  suppress number;
}
```

Hierarchy Level

[edit interfaces *interface--name*],
[edit interfaces *interface--range*]

Release Information

Statement introduced in Junos OS Release 14.1 for PTX Series Packet Transport Routers and T Series Core Routers.
Statement introduced in Junos OS Release 14.2 for MX960, MX480, MX240, and MX80 3D Universal Edge Routers and M10i Multiservice Edge Routers.

Description

Limit the number of advertisements of the up and down transitions (flapping) on an interface. Each time a transition occurs, the interface state is changed, which generates an advertisement to the upper-level routing protocols. Damping helps reduce the number of these advertisements. Every time an interface goes down, a penalty is added to the interface penalty counter. Penalty added on every interface flap is 1000.

If at some point the accumulated penalty exceeds the suppress level **max-suppress**, the interface is placed in the suppress state, and further interface state up and down transitions are not reported to the upper-level protocols.

Options

enable—Enable damping on a per-interface basis. If damping is enabled on an interface, it is suppressed during interface flaps that match the configuration settings.

Default: Disabled

half-life ***seconds***—Decay half-life. ***seconds*** is the interval after which the accumulated interface penalty counter is reduced by half if the interface remains stable.

NOTE: **For the half-life, configure a value that is less than the max-suppress value. If you do not, the configuration is rejected.**

Range: 1 through 30

Default: 5

max-suppress ***seconds***—Maximum hold-down time. ***seconds*** is the maximum time that an interface can be suppressed no matter how unstable the interface has been.

Copyright © 2018, Juniper Networks, Inc.

NOTE: For max-suppress, configure a value that is greater than the half-life. If you do not, the configuration is rejected.

Range: 1 through 20,000

Default: 20

reuse *number*—Reuse threshold. When the accumulated interface penalty counter falls below *number*, the interface is no longer suppressed.

Range: 1 through 20,000

Default: 1000

suppress *number*—Cutoff (suppression) threshold. When the accumulated interface penalty counter exceeds *number*, the interface is suppressed.

Range: 1 through 20,000

Default: 2000

Required Privilege Level

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

Related Documentation

- Physical Interface Damping Overview on page 159
- Damping Shorter Physical Interface Transitions on page 165
- Damping Longer Physical Interface Transitions on page 166
- **show interfaces extensive on page 1555**
- **hold-time on page 644**

Copyright © 2018, Juniper Networks, Inc.

data-channel

Syntax

```
data-channel {
  vlan number;
}
```

Hierarchy Level [edit protocols protection-group ethernet-ring *ring-name*]

Release Information Statement introduced in Junos OS Release 10.2.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

Description For Ethernet ring protection, configure a data channel to define a set of VLAN IDs that belong to a ring instance.

VLANs specified in the data channel use the same topology used by the ERPS PDU in the control channel. Therefore, if a ring interface is blocked in the control channel, all traffic in the data channel is also blocked on that interface.

Options **vlan *number***—Specify (by VLAN ID) one or more VLANs that belong to a ring instance.

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

Related Documentation

- *Ethernet Ring Protection Using Ring Instances for Load Balancing*
- *Example: Configuring Load Balancing Within Ethernet Ring Protection for MX Series Routers*
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

Copyright © 2018, Juniper Networks, Inc.

data-tlv-size

Syntax data-tlv-size *size*;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id* remote-mep *remote-mep-id* sla-iterator-profile *profile-name*]

Release Information Statement introduced in Junos OS Release 11.1.

Description Configure the size of the data TLV portion of the Y.1731 data frame.

Options *size*—Size of the data TLV portion of the Y.1731 data frame.

NOTE: This option is applicable only for two-way delay measurement.

Range: 1 through 1400 bytes

Default: 1

Required Privilege Level Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- sla-iterator-profile on page 969
- *Configuring a Remote MEP with an Iterator Profile*

 Copyright © 2018, Juniper Networks, Inc.

dcd

Syntax dcd (ignore | normal | require);

Hierarchy Level [edit interfaces *interface-name* serial-options dce-options],
[edit interfaces *interface-name* serial-options dte-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For EIA-530 and V.35 interfaces only, configure the from-DCE signal, data-carrier-detect (DCD).

Options **ignore**—The from-DCE signal is ignored.

normal—Normal DCD signal handling as defined by the TIA/EIA Standard 530.

require—The from-DCE signal must be asserted.

Default: normal

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

Related Documentation

- Configuring the Serial Signal Handling on page 344

dcd-polarity

Syntax dcd-polarity (negative | positive);

Hierarchy Level [edit interfaces *interface-name* serial-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure DCD signal polarity.

Options **positive**—Positive signal polarity.

negative—Negative signal polarity.

Default: positive

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

Related Documentation

- Configuring Serial Signal Polarities on page 347

Copyright © 2018, Juniper Networks, Inc.

dce

Syntax dce;

Hierarchy Level [edit interfaces *interface-name*],
[edit interfaces *interface-name* serial-options clocking-mode]

Release Information Statement introduced before Junos OS Release 7.4.

Description For Frame Relay only, respond to status enquiry message keepalives.

When you configure the router to be a DCE, keepalives are disabled by default.

Default The router operates in DTE mode.

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

Related Documentation

- Configuring the Router as a DCE with Frame Relay Encapsulation on page 154

deactivation-delay

Syntax deactivation-delay *seconds*;

Hierarchy Level [edit interfaces dl*n* unit *logical-unit-number* dialer-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description On J Series Services Routers with ISDN interfaces, configure the ISDN deactivation delay. Used only for dialer backup and dialer watch cases.

Options *seconds*—Interval before the backup interface is deactivated after the primary interface has comes up.

Range: 1 through 4,294,967,295 seconds

Default: 0 (zero)

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

Related Documentation

- *Junos OS Interfaces and Routing Configuration Guide*

 Copyright © 2018, Juniper Networks, Inc.

dce-options

Syntax

```
dce-options {
  control-signal (assert | de-assert | normal);
  cts (ignore | normal | require);
  dcd (ignore | normal | require);
  dsr (ignore | normal | require);
  dtr signal-handling-option;
  ignore-all;
  indication (ignore | normal | require);
  rts (assert | de-assert | normal);
  tm (ignore | normal | require);
}
```

Hierarchy Level

[edit interfaces *interface-name* serial-options]

Release Information

Statement introduced in Junos OS Release 8.3.
Statement previously known as **control-leads**.

Description

For J Series Services Routers, configure the serial interface signal characteristics.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- Configuring the Serial Signal Handling on page 344

default-actions

Syntax
```
default-actions {
  interface-down;
}
```

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management action-profile *profile-name*]

Release Information Statement introduced in Junos OS Release 8.4.

Description Define the action to be taken when connectivity to the remote MEP is lost.

Default If no action is configured, no action is taken.

Options **interface-down**—When a remote MEP connectivity failure is detected, bring the interface down.

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

Related Documentation
- *Configuring a CFM Action Profile to Specify CFM Actions for CFM Events*

 Copyright © 2018, Juniper Networks, Inc.

default-chap-secret

Syntax default-chap-secret *name*;

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

Release Information Statement introduced in Junos OS Release 8.0.

Description Define the default CHAP secret to be used when no matching CHAP access profile exists.

For ATM2 IQ interfaces only, you can configure a default CHAP secret on the logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:

- **atm-ppp-llc**—PPP over AAL5 LLC encapsulation.
- **atm-ppp-vc-mux**—PPP over AAL5 multiplex encapsulation.

Default If you do not include the **default-chap-secret** statement in the configuration, and an interface receives a CHAP challenge or response from a peer that is not in the applied access profile, the link is immediately dropped.

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

Related Documentation

- *Configuring a Default CHAP Secret*
- access-profile on page 412

Copyright © 2018, Juniper Networks, Inc.

default-pap-password

Syntax default-pap-password *password*;

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

Release Information Statement introduced in Junos OS Release 8.3.

Description For PAP authentication, the default PAP password.

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

Related Documentation

- *Configuring a Default PAP Password*
- access-profile on page 412

delimiter

Syntax delimiter *delimiter-character*;

Hierarchy Level [edit interfaces *interface-name* auto-configure vlan-ranges authentication username-include],
[edit interfaces *interface-name* auto-configure stacked-vlan-ranges authentication username-include]

Release Information Statement introduced in Junos OS Release 10.0.

Description Specify the character used as the delimiter between the concatenated components of the username. You cannot use the semicolon (;) as a delimiter.

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

Related Documentation

- *Configuring VLAN Interface Username Information for AAA Authentication*

 Copyright © 2018, Juniper Networks, Inc.

demux-destination (Underlying Interface)

Syntax demux-destination *family*;

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

Release Information Statement introduced in Junos OS Release 9.0.
Support for aggregated Ethernet added in Junos OS Release 9.4.

Description Configure the logical demultiplexing (demux) destination family type on the IP demux underlying interface.

NOTE: The IP demux interface feature currently supports only Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet underlying interfaces.

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

Related Documentation

- Configuring an IP Demultiplexing Interface on page 316
- Configuring a VLAN Demultiplexing Interface on page 320

demux-destination (Demux Interface)

Syntax

demux-destination {
  destination-prefix;
}

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* family *family*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

Release Information

Statement introduced in Junos OS Release 9.0.
Support for aggregated Ethernet added in Junos OS Release 9.4.

Description

Configure one or more logical demultiplexing (demux) destination prefixes. The prefixes are matched against the destination address of packets that the underlying interface receives. When a match occurs, the packet is processed as if it was received on the demux interface.

Required Privilege Level

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

Related Documentation

- Configuring an IP Demultiplexing Interface on page 316
- Configuring a VLAN Demultiplexing Interface on page 320

 Copyright © 2018, Juniper Networks, Inc.

demux-options (Static Interface)

Syntax

```
demux–options {
  underlying-interface interface-name
}
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 9.0.

Description

Configure logical demultiplexing (demux) interface options.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- Configuring an IP Demultiplexing Interface on page 316
- Configuring a VLAN Demultiplexing Interface on page 320

demux-source (Demux Interface)

Syntax

```
demux-source {
  source-prefix;
}
```

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* family *family*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

Release Information

Statement introduced in Junos OS Release 9.0.
Support for aggregated Ethernet added in Junos OS Release 9.4.

Description

Configure one or more logical demultiplexing (demux) source prefixes. The prefixes are matched against the source address of packets that the underlying interface receives. When a match occurs, the packet is processed as if it was received on the demux interface.

Required Privilege Level

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

Related Documentation

- Configuring an IP Demultiplexing Interface on page 316
- Configuring a VLAN Demultiplexing Interface on page 320

Copyright © 2018, Juniper Networks, Inc.

demux-source (Underlying Interface)

Syntax demux-source *family*;

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

Release Information Statement introduced in Junos OS Release 9.0.
Support for aggregated Ethernet added in Junos OS Release 9.4.

Description Configure the logical demultiplexing (demux) source family type on the IP demux underlying interface.

NOTE: The IP demux interface feature currently supports only Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet underlying interfaces.

Options *family*—Protocol family:

- **inet**—Internet Protocol version 4 suite
- **inet6**—Internet Protocol version 6 suite

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

Related Documentation

- Configuring an IP Demultiplexing Interface on page 316
- Configuring a VLAN Demultiplexing Interface on page 320

 Copyright © 2018, Juniper Networks, Inc.

demux0 (Static Interface)

Syntax

```
demux0 {
  unit logical-unit-number {
    demux-options {
      underlying-interface interface-name
    }
    family family {
      access-concentrator name;
      {
        destination-prefix;
      }
      direct-connect;
      duplicate-protection;
      dynamic-profile profile-name;
      {
        source-prefix;
      }
      max-sessions number;
      service-name-table table-name
      targeted-distribution;
      unnumbered-address interface-name <preferred-source-address address>;
    }
    vlan-id number;
    vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
  }
}
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 9.0.

Description

Configure the logical demultiplexing (demux) interface.

Logical IP demux interfaces do not support IPv4 and IPv6 dual stack.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- Configuring an IP Demultiplexing Interface on page 316
- Configuring a VLAN Demultiplexing Interface on page 320

demux0 (Dynamic Interface)

Syntax

```
demux0 {
  unit logical-unit-number {
    demux-options {
      underlying-interface interface-name
    }
    family family {
      access-concentrator name;
      address address;
      demux-source {
        source-prefix;
      }
      direct-connect;
      duplicate-protection;
      dynamic-profile profile-name;
      filter {
        input filter-name;
        output filter-name;
      }
      mac-validate (loose | strict):
      max-sessions number;
      max-sessions-vsa-ignore;
      rpf-check {
        fail-filter filter-name;
        mode loose;
      }
      service-name-table table-name
      short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
        maximum-seconds>;
      unnumbered-address interface-name <preferred-source-address address>;
    }
    filter {
      input filter-name;
      output filter-name;
    }
    vlan-id number;
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces]

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Configure the logical demultiplexing (demux) interface in a dynamic profile.

Logical IP demux interfaces do not support IPv4 and IPv6 dual stack.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Copyright © 2018, Juniper Networks, Inc.

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

Related Documentation
- *Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles*
- Demultiplexing Interface Overview on page 313

Copyright © 2018, Juniper Networks, Inc.

description (Interfaces)

Syntax description *text*;

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

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Provide a textual description of the interface or the logical unit. Any descriptive text you include is displayed in the output of the **show interfaces** commands, and is also exposed in the **ifAlias** Management Information Base (MIB) object. It has no effect on the operation of the interface on the router or switch.

The textual description can also be included in the extended DHCP relay option 82 Agent Circuit ID suboption.

Options *text*—Text to describe the interface. If the text includes spaces, enclose the entire text in quotation marks.

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

Related Documentation

- Configuring Interface Description on page 113
- Adding a Logical Unit Description to the Configuration on page 188
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit and 10-Gigabit Ethernet Interfaces for OCX Series Switches*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure) for EX Series Switches with ELS support*
- *Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX and QFX Series Switches*
- *Using DHCP Relay Agent Option 82 Information*
- *Junos OS Network Interfaces Library for Routing Devices*
- *Example: Connecting Access Switches with ELS Support to a Distribution Switch with ELS Support*

 Copyright © 2018, Juniper Networks, Inc.

destination (IPCP)

Syntax
destination *address* destination-profile *profile-name*;

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

Release Information
Statement introduced before Junos OS Release 7.4.

Description
For unnumbered interfaces with PPP encapsulation, specify the IP address of the remote interface.

Options
address—IP address of the remote interface.

The remaining statement is explained separately. See CLI Explorer.

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

Related Documentation
- Configuring IPCP Options for Interfaces with PPP Encapsulation on page 216
- **address on page 424**
- **negotiate-address on page 805**
- *Junos OS Administration Library*

destination (Routing Instance)

Syntax destination *routing-instance-name*;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the destination routing instance that points to the routing table containing the tunnel destination address.

Default The default Internet routing table **inet.0**.

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

 Copyright © 2018, Juniper Networks, Inc.

destination (Tunnels)

Syntax

destination *address*;

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*],
[edit interfaces *interface-name* unit *logical-unit-number* family inet unnumbered-address *interface-name*],
[edit interfaces *interface-name* unit *logical-unit-number* tunnel],
[edit logical-systems *logical-system-name* interfaces interface-name unit *logical-unit-number* family inet address *address*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet unnumbered-address *interface-name*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* tunnel]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

For encrypted, PPP-encapsulated, and tunnel interfaces, specify the remote address of the connection.

Options

address—Address of the remote side of the connection.

Required Privilege Level

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

Related Documentation

- Configuring the Interface Address on page 207
- *Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)*
- *Junos OS Services Interfaces Library for Routing Devices*

destination-class-usage

Syntax destination-class-usage;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Enable packet counters on an interface that count packets that arrive from specific customers and are destined for specific prefixes on the provider core router.

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

Related Documentation

- Enabling Source Class and Destination Class Usage on page 262
- accounting on page 413
- source-class-usage on page 977

 Copyright © 2018, Juniper Networks, Inc.

destination-profile

Syntax

destination-profile *name*;

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*],
[edit interfaces *interface-name* unit *logical-unit-number* family inet unnumbered-address *interface-name* destination *address*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet unnumbered-address *interface-name* destination *address*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description

For interfaces with PPP encapsulation, assign PPP properties to the remote destination end. You define the profile at the **[edit access group-profile *name* ppp]** hierarchy level.

Options

name—Profile name defined at the **[edit access group-profile *name* ppp]** hierarchy level.

Required Privilege Level

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

Related Documentation

- Configuring IPCP Options for Interfaces with PPP Encapsulation on page 216
- **destination (IPCP) on page 527**
- *Junos OS Administration Library*

dial-options

Syntax

```
dial-options {
  ipsec-interface-id name;
  l2tp-interface-id name;
  (shared | dedicated);
}
```

Hierarchy Level

[edit interfaces sp-*fpc/pic/port* unit *logical-unit-number*],
[edit interfaces si-*fpc/pic/port* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces sp-*fpc/pic/port* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces si-*fpc/pic/port* unit *logical-unit-number*]

Release Information

Statement introduced before Junos OS Release 7.4.
The **[edit ...si-...]** hierarchy levels introduced in Junos OS Release 11.4.

Description

Specify the options for configuring logical interfaces for group and user sessions in L2TP or IPsec dynamic endpoint tunneling.

Options

dedicated—(LNS on M Series routers and MX Series routers only) Specify that a logical interface can host only one session at a time.

ipsec-interface-id *name*—(M Series routers only) Interface identifier for group of dynamic peers. This identifier must be replicated at the **[edit access profile *name* client * ike]** hierarchy level.

l2tp-interface-id *name*—Interface identifier that must be replicated at the **[edit access profile *name*]** hierarchy level.

shared—(LNS on M Series routers only) Specify that a logical interface can host multiple (shared) sessions at a time.

Required Privilege Level

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

Related Documentation

- *Configuring the Identifier for Logical Interfaces that Provide L2TP Services*
- *Configuring Dynamic Endpoints for IPsec Tunnels*
- *Configuring Options for the LNS Inline Services Logical Interface*

 Copyright © 2018, Juniper Networks, Inc.

dial-string

Syntax dial-string [*dial-string-numbers*];

Hierarchy Level [edit interfaces br-*pim/0/port* unit *logical-unit-number* dialer-options],
[edit logical-systems *logical-system-name* interfaces br-*pim/0/port* unit *logical-unit-number* dialer-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description On J Series Services Routers with ISDN interfaces, specify one or more ISDN dial strings used to reach a destination subnetwork.

Options ***dial-string-numbers***—One or more strings of numbers to call.

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

Related Documentation None

dialer

Syntax dialer *filter-name*;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Apply a dialer filter to an interface. To create the dialer filter, include the **dialer-filter** statement at the **[edit firewall filter family *family*]** hierarchy level.

Options ***filter-name***—Dialer filter name.

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

Related Documentation
- *Junos OS Interfaces and Routing Configuration Guide*

dialer-options

Syntax

```
dialer-options {
  activation-delay seconds;
  callback;
  callback-wait-period time;
  deactivation-delay seconds;
  dial-string [ dial-string-numbers ];
  idle-timeout seconds;
  incoming-map {
    caller caller-number | accept-all;
    initial-route-check seconds;
    load-interval seconds;
    load-threshold percent;
    pool pool-name;
    redial-delay time;
    watch-list {
      [ routes ];
    }
  }
}
```

Hierarchy Level

[edit interfaces umd0],
[edit interfaces dl*n* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces dl*n* unit *logical-unit-number*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the dialer options for configuring logical interfaces for group and user sessions.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2018, Juniper Networks, Inc.

dialin

Syntax dialin (console | routable);

Hierarchy Level [edit interfaces umd0 modem-options]

Release Information Statement introduced in Junos OS Release 8.5.

Description For J Series Services Routers, configure a USB modem port to act as a dial-in console or WAN backup port.

Options **console**—Configure the USB modem port to operate as a dial-in console for management.

routable—Configure the USB modem port to operate as a dial-in WAN backup interface.

Default: console

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

Related Documentation

- Specifying a USB Modem Interface on J Series Routers on page 350

direction

Syntax direction (up | down);

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id*]

Release Information Statement introduced in Junos OS Release 8.4.

Description Configure the direction of the MEP.

Options **up**—An UP MEP CCM is transmitted out of every logical interface which is part of the same bridging or vpls instance except for the interface configured on this MEP.

NOTE: The up direction for MEP is not supported on T Series routers.

down—Down MEP CCMs are transmitted only out the interface configured on this MEP.

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

Related Documentation

- *Configuring a MEP to Generate and Respond to CFM Protocol Messages*
- *IEEE 802.1ag OAM Connectivity Fault Management Overview*

Copyright © 2018, Juniper Networks, Inc.

disable (Interface)

Syntax

disable;

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description

Disable a physical or a logical interface, effectively unconfiguring it.

CAUTION:

- Dynamic subscribers and logical interfaces use physical interfaces for connection to the network. The Junos OS allows you to set the interface to disable and commit the change while dynamic subscribers and logical interfaces are still active. This action results in the loss of all subscriber connections on the interface. Use care when disabling interfaces.
- If aggregated SONET links are configured between a T1600 router and a T4000 router, interface traffic is disrupted when you disable the physical interface configured on the T1600 router. If you want to remove the interface, we recommend that you deactivate the interface instead of disabling it.

NOTE:

- When you use the **disable** statement at the **[edit interfaces]** hierarchy level, depending on the PIC type, the interface might or might not turn off the laser. Older PIC transceivers do not support turning off the laser, but newer Gigabit Ethernet (GE) PICs with SFP and XFP transceivers and ATM MIC with SFP do support it and the laser will be turned off when the interface is disabled. If the ATM MIC with SFP is part of an APS group, then the laser will not be turned off when you use the **disable** statement at the **[edit interfaces]** hierarchy level..
- When you disable or deactivate an interface, then all the references made to the deactivated interface must be removed from the routing instance.
- For abstracted fabric interfaces, the **disable** command disables AF interface on the local GNF only.

Copyright © 2018, Juniper Networks, Inc.

WARNING: **Do not stare into the laser beam or view it directly with optical instruments even if the interface has been disabled.**

Required Privilege Level

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

Related Documentation

- Disabling a Physical Interface on page 174
- Disabling a Logical Interface on page 201

disable (Link Protection)

Syntax

disable;

Hierarchy Level

[edit interfaces ae*X* aggregated-ether-options lacp link-protection]

Release Information

Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 11.4 for EX Series switches.
Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.

Description

Disable LACP link protection on the interface.

Required Privilege Level

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

Related Documentation

- *Configuring LACP for Aggregated Ethernet Interfaces*
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*

Copyright © 2018, Juniper Networks, Inc.

disable-mlppp-inner-ppp-pfc

Syntax disable-mlppp-inner-ppp-pfc;

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

Release Information Statement introduced in Junos OS Release 8.2.

Description For MLPPP interfaces only, disable compression of the inner PPP header in the MLPPP payload. By default, compression is enabled.

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

Related Documentation
- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2018, Juniper Networks, Inc.

dlci

Syntax dlci *dlci-identifier*;

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

Release Information Statement introduced before Junos OS Release 7.4.
Starting with Junos OS Release 18.2R1, the SRX Series devices support frame-relay encapsulation and adds DLCI information to the given frame.

Description For Frame Relay and Multilink Frame Relay (MLFR) user-to-network interface (UNI) network-to-network interface (NNI) encapsulation only, and for link services, voice services and point-to-point interfaces only, configure the data-link connection identifier (DLCI) for a permanent virtual circuit (PVC) or an switched virtual circuit (SVC). The DLCI setups a frame-relay PVC to form a L2 point-to-point connection. This is used for peering different LT IFL pairs.

To configure a DLCI for a point-to-multipoint interface, use the **multipoint-destination** statement to specify the DLCI.

Options ***dlci-identifier***—Data-link connection identifier.

Range: 16 through 1022.

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

Related Documentation

- *Data-Link Connection Identifiers on Channelized Interfaces*
- *Configuring Frame Relay DLCIs*
- *Junos OS Services Interfaces Library for Routing Devices*
- encapsulation (Logical Interface) on page 573
- multicast-dlci on page 790
- multipoint-destination on page 794

 Copyright © 2018, Juniper Networks, Inc.

do-not-fragment

Syntax do-not-fragment;

Hierarchy Level [edit interfaces gr-*fpc/pic/port* unit *logical-unit-number* tunnel],
[edit logical-systems *logical-system-name* interfaces gr-*fpc/pic/port* unit *logical-unit-number* tunnel]

Release Information Statement introduced in Junos OS Release 9.2.
Statement introduced in Junos OS Release 15.1X53-D10 for QFX10000 switches.

Description For a generic routing encapsulation (GRE) tunnel, disable fragmentation of GRE-encapsulated packets. This sets the do-not-fragment (DF) bit in the outer IP header of the GRE-encapsulated packets so that they do not get fragmented anywhere in the path. When the size of a GRE-encapsulated packet is greater than the MTU of a link that the packet passes through, the GRE-encapsulated packet is dropped.

Default By default, fragmentation of GRE-encapsulated packets is disabled.

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

Related Documentation

- allow-fragmentation on page 436
- *reassemble-packets*
- *Enabling Fragmentation and Reassembly on Packets After GRE-Encapsulation*
- *Junos OS Services Interfaces Library for Routing Devices*

domain-name

Syntax domain-name *domain-name-string*;

Hierarchy Level [edit interfaces *interface-name* auto-configure vlan-ranges authentication username-include],
[edit interfaces *interface-name* auto-configure stacked-vlan-ranges authentication username-include]

Release Information Statement introduced in Junos OS Release 10.0.

Description Specify the domain name that is concatenated with the username during the subscriber authentication process.

Options ***domain-name-string***—The domain name formatted string.

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

Related Documentation
- *Configuring VLAN Interface Username Information for AAA Authentication*

 Copyright © 2018, Juniper Networks, Inc.

dot1x

Syntax

```
dot1x {
  authenticator {
    authentication-profile-name access-profile-name;
    interface interface-id {
      maximum-requests integer;
      quiet-period seconds;
      reauthentication (disable | interval seconds);
      retries integer;
      server-timeout seconds;
      supplicant (single);
      supplicant-timeout seconds;
      transmit-period seconds;
    }
  }
}
```

Hierarchy Level [edit protocols]

Release Information Statement introduced in Junos OS Release 9.3.

Description For the MX Series only, specifies settings for using 802.1x Port-Based Network Access Control.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *IEEE 802.1x Port-Based Network Access Control Overview*
- authenticator on page 448
- authentication-profile-name on page 447
- interface (IEEE 802.1x) on page 682

down-count

Syntax down-count *cells*;

Hierarchy Level [edit interfaces *interface-name* atm-options vpi *vpi-identifier* oam-liveness],
[edit interfaces *interface-name* unit *logical-unit-number* oam-liveness],
[edit interfaces *interface-name* unit *logical-unit-number* family *family* address *address* multipoint-destination *address* oam-liveness],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* oam-liveness],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family* address *address* multipoint-destination *address* oam-liveness]

Release Information Statement introduced before Junos OS Release 7.4.

Description For ATM encapsulation only, configure Operation, Administration, and Maintenance (OAM) F5 loopback cell count thresholds. This feature is not supported on ATM-over-SHDSL interfaces.

For ATM2 IQ PICs only, configure OAM F4 loopback cell count thresholds at the **[edit interfaces *interface-name* atm-options vpi *vpi-identifier*]** hierarchy level.

Options *cells*—Minimum number of consecutive OAM F4 or F5 loopback cells lost before a VC is declared down.

Range: 1 through 255

Default: 5 cells

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

Related Documentation

- *Configuring the ATM OAM F5 Loopback Cell Threshold*

 Copyright © 2018, Juniper Networks, Inc.

drop (PPPoE Service Name Tables)

Syntax

drop;

Hierarchy Level

[edit protocols pppoe service-name-tables *table-name* service *service-name*],
[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]

Release Information

Statement introduced in Junos OS Release 10.0.
Support at **[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]** hierarchy level introduced in Junos OS Release 10.2.

Description

Direct the router to drop (ignore) a PPPoE Active Discovery Initiation (PADI) control packet received from a PPPoE client that contains the specified service name tag or agent circuit identifier/agent remote identifier (ACI/ARI) information. This action effectively denies the client's request to provide the specified service, or to accept requests from the subscriber or subscribers represented by the ACI/ARI information.

Required Privilege Level

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

Related Documentation

- *Configuring PPPoE Service Name Tables*

Copyright © 2018, Juniper Networks, Inc.

drop-timeout

Syntax drop-timeout *milliseconds*;

Hierarchy Level [edit interfaces *interface-name* mlfr-uni-nni-bundle-options],
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services, multilink, and voice services interfaces only, configure the drop timeout period, in milliseconds.

Options ***milliseconds***—Drop timeout period.

Range: 0 through 2000 milliseconds

Default: 0 ms (disabled)

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2018, Juniper Networks, Inc.

ds0-options

Syntax

```
ds0-options {
  bert-algorithm algorithm;
  bert-error-rate rate;
  bert-period seconds;
  byte-encoding (nx56 | nx64);
  fcs (16 | 32);
  idle-cycle-flag (flags | ones);
  invert-data;
  loopback payload;
  start-end-flag (filler | shared);
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure DS0-specific physical interface properties.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *Configuring Channelized DS3-to-DS0 Interfaces*

Copyright © 2018, Juniper Networks, Inc.

dsl-options

Syntax

```
dsl-options {
  loopback local;
  operating-mode mode;
}
```

Hierarchy Level [edit interfaces at-*fpc/pic/port*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For J Series Services Routers only, modify the properties of the digital subscriber line for an ATM interface.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *ATM-over-ADSL Overview*
- *Junos OS Interfaces and Routing Configuration Guide*

 Copyright © 2018, Juniper Networks, Inc.

dsr

Syntax dsr (ignore | normal | require);

Hierarchy Level [edit interfaces *interface-name* serial-options dce-options],
[edit interfaces *interface-name* serial-options dte-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For EIA-530 and V.35 interfaces only, configure the from-DCE signal, data-set-ready (DSR).

Options **ignore**—The from-DCE signal is ignored.

normal—Normal DSR signal handling as defined by the TIA/EIA Standard 530.

require—The from-DCE signal must be asserted.

Default: normal

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

Related Documentation

- Configuring the Serial Signal Handling on page 344

dsr-polarity

Syntax dsr-polarity (negative | positive);

Hierarchy Level [edit interfaces *interface-name* serial-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure DSR signal polarity.

Options **positive**—Positive signal polarity.

negative—Negative signal polarity.

Default: positive

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

Related Documentation

- Configuring Serial Signal Polarities on page 347

dte-options

Syntax

```
dte-options {
  control-signal (assert | de-assert | normal);
  cts (ignore | normal | require);
  dcd (ignore | normal | require);
  dsr (ignore | normal | require);
  dtr signal-handling-option;
  ignore-all;
  indication (ignore | normal | require);
  rts (assert | de-assert | normal);
  tm (ignore | normal | require);
}
```

Hierarchy Level

[edit interfaces *interface-name* serial-options]

Release Information

Statement introduced in Junos OS Release 8.3.
Statement previously known as **control-leads**.

Description

For M Series and T Series routers, configure the serial interface signal characteristics.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- Configuring the Serial Signal Handling on page 344

Copyright © 2018, Juniper Networks, Inc.

dtr

Syntax

dtr *signal-handling-option*;

Hierarchy Level

[edit interfaces *interface-name* serial-options dce-options],
[edit interfaces *interface-name* serial-options dte-options]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

For EIA-530 and V.35 interfaces only, configure the to-DCE signal, data-transmit-ready (DTR).

Options

signal-handling-option—Signal handling for the DTR signal. The signal handling can be one of the following:

assert—The to-DCE signal must be asserted.

auto-synchronize—Normal DTR signal with automatic synchronization. This statement has two substatements:

duration *milliseconds*—Pulse duration of resynchronization.

Range: 1 through 1000 milliseconds

Default: 1000 milliseconds

interval *seconds*—Offset interval for resynchronization.

Range: 1 through 31 seconds

Default: 15 seconds

de-assert—The to-DCE signal must be deasserted.

normal—Normal DTR signal handling as defined by the TIA/EIA Standard 530.

Default: **normal**

Required Privilege Level

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

Related Documentation

- Configuring the Serial Signal Handling on page 344

dtr-circuit

Syntax dtr-circuit (balanced | unbalanced);

Hierarchy Level [edit interfaces *interface-name* serial-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For EIA-530 and V.35 interfaces only, configure a DTR circuit.

Options **balanced**—Balanced DTR signal.

unbalanced—Unbalanced DTR signal.

Default: balanced

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

Related Documentation

- Configuring the Serial DTR Circuit on page 347

dtr-polarity

Syntax dtr-polarity (negative | positive);

Hierarchy Level [edit interfaces *interface-name* serial-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure DTR signal polarity.

Options **positive**—Positive signal polarity.

negative—Negative signal polarity.

Default: positive

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

Related Documentation

- Configuring Serial Signal Polarities on page 347

 Copyright © 2018, Juniper Networks, Inc.

dump-on-flow-control

Syntax dump-on-flow-control;

Hierarchy Level [edit interfaces *interface-name* multiservice-options]

Release Information Statement introduced in Junos OS Release 9.5.

Description This option supports high availability functionality and can be used with various service interfaces, including **rsp**, **rms**, **lsq**, and **rlsq**.

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

Related Documentation

- Configuring Multiservice Physical Interface Properties on page 170
- *Junos OS Services Interfaces Library for Routing Devices*
- passive-monitor-mode on page 857

Copyright © 2018, Juniper Networks, Inc.

dynamic-call-admission-control

Syntax

```
dynamic-call-admission-control {
  activation-priority priority;
  bearer-bandwidth-limit kilobits-per-second;
}
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 8.2.

Description

(J4350 and J6350 Services Routers supporting voice over IP with the TGM550 media gateway module) For Fast Ethernet and Gigabit Ethernet interfaces, ISDN BRI interfaces, and serial interfaces with PPP or Frame Relay encapsulation, configure dynamic call admission control (CAC). Dynamic CAC provides enhanced control over WAN bandwidth. When dynamic CAC is configured on an interface responsible for providing call bandwidth, the TGM550 informs the Media Gateway Controller (MGC) of the bandwidth limit available for voice packets on the interface and requests the MGC to block new calls when the bandwidth is exhausted.

Dynamic CAC must be configured on each Services Router interface responsible for providing call bandwidth.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Configuring Dynamic Call Admission Control*
- *Junos OS Services Interfaces Library for Routing Devices*
- *Junos OS Interfaces and Routing Configuration Guide*

Copyright © 2018, Juniper Networks, Inc.

dynamic-profile (PPP)

Syntax

dynamic-profile *profile-name*;

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 9.5.
Support for MLPPP on LSQ interfaces introduced in Junos OS Release 10.2.

Description

Specify the dynamic profile that is attached to the interface. On the MX Series routers, this statement is supported on PPPoE interfaces only.

Required Privilege Level

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

Related Documentation

- *Dynamic Profiles Overview*
- *Configuring a Basic Dynamic Profile*
- *Attaching Dynamic Profiles to Static PPP Subscriber Interfaces*
- *Attaching Dynamic Profiles to MLPPP Bundles*
- For hardware requirements, see *Hardware Requirements for PPP Subscriber Services on Non-Ethernet Interfaces*

Copyright © 2018, Juniper Networks, Inc.

dynamic-profile (PPPoE Service Name Tables)

Syntax dynamic-profile *profile-name*;

Hierarchy Level [edit protocols pppoe service-name-tables *table-name* service *service-name*],
[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify a dynamic profile to instantiate a dynamic PPPoE interface. You can associate a dynamic profile with a named service entry, **empty** service entry, or **any** service entry configured in a PPPoE service name table, or with an agent circuit identifier/agent remote identifier (ACI/ARI) pair defined for these services.

The dynamic profile associated with a service entry in a PPPoE service name table overrides the dynamic profile associated with the PPPoE underlying interface on which the dynamic PPPoE interface is created.

If you include the **dynamic-profile** statement at the **[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]** hierarchy level, you cannot also include the **static-interface** statement at this level. The **dynamic-profile** and **static-interface** statements are mutually exclusive for ACI/ARI pair configurations.

Options ***profile-name***—Name of the dynamic profile that the router uses to instantiate a dynamic PPPoE interface.

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

Related Documentation

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

 Copyright © 2018, Juniper Networks, Inc.

dynamic-profile (Stacked VLAN)

Syntax

```
dynamic-profile profile-name {
  accept (any | dhcp-v4 |dhcp-v6| inet | inet6 | pppoe);
  access-profile vlan-dynamic-profile-name;
  ranges (any | low-tag–high-tag),(any | low-tag–high-tag);
}
```

Hierarchy Level [edit interfaces *interface-name* auto-configure stacked-vlan-ranges]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure a dynamic profile for use when configuring dynamic stacked VLANs.

Options ***profile-name***—Name of the dynamic profile that you want to use when configuring dynamic stacked VLANs.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

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

Related Documentation

- *Dynamic Profiles Overview*
- *Configuring a Basic Dynamic Profile*
- *Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs*

Copyright © 2018, Juniper Networks, Inc.

dynamic-profile (VLAN)

Syntax

```
dynamic-profile profile-name {
  accept (any | dhcp-v4 |dhcp-v6| inet | inet6 | pppoe);
  accept-out-of-band protocol;
  access-profilevlan-dynamic-profile-name;
  ranges (any | low-tag)–(any | high-tag);
}
```

Hierarchy Level

[edit interfaces *interface-name* auto-configure vlan-ranges]

Release Information

Statement introduced in Junos OS Release 9.5.

Description

Configure a dynamic profile for use when configuring dynamic VLANs.

Options

profile-name—Name of the dynamic profile that you want to use when configuring dynamic VLANs.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

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

Related Documentation

- *Dynamic Profiles Overview*
- *Configuring a Basic Dynamic Profile*
- *Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs*

 Copyright © 2018, Juniper Networks, Inc.

dynamic-profiles

Syntax

```
dynamic-profiles {
  profile-name {
    class-of-service {
      interfaces {
        interface-name ;
        }
        unit logical-unit-number {
          classifiers {
            type (classifier-name | default);
          }
          output-traffic-control-profile (profile-name | $junos-cos-traffic-control-profile);
          report-ingress-shaping-rate bps;
          rewrite-rules {
            dscp (rewrite-name | default);
            dscp-ipv6 (rewrite-name | default);
            ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
            inet-precedence (rewrite-name | default);
            }
          }
        }
      }
      scheduler-maps {
        map-name {
          forwarding-class class-name scheduler scheduler-name;
        }
      }
      schedulers {
        (scheduler-name) {
          buffer-size (seconds | percent percentage | remainder | temporal microseconds);
          drop-profile-map loss-priority (any | low | medium-low | medium-high | high)
            protocol (any | non-tcp | tcp) drop-profile profile-name;
          excess-priority (low | high | $junos-cos-scheduler-excess-priority);
          excess-rate (percent percentage | percent $junos-cos-scheduler-excess-rate);
          overhead-accounting (shaping-mode) <bytes (byte-value>;
          priority priority-level;
          shaping-rate (rate | predefined-variable);
          transmit-rate (percent percentage | rate | remainder) <exact | rate-limit>;
        }
      }
      traffic-control-profiles profile-name {
        delay-buffer-rate (percent percentage | rate | $junos-cos-delay-buffer-rate);
        excess-rate (percent percentage | proportion value | percent $junos-cos-excess-rate);
        guaranteed-rate (percent percentage | rate | $junos-cos-guaranteed-rate);
        overhead-accounting (shaping-mode) <bytes (byte-value>;
        scheduler-map map-name;
        shaping-rate (rate | predefined-variable);
      }
    }
    firewall {
      family family {
        fast-update-filter filter-name {
          interface-specific;
```

Copyright © 2018, Juniper Networks, Inc.

```
        match-order [match-order];
        term term-name {
          from {
            match-conditions;
          }
          then {
            action;
            action-modifiers;
          }
          only-at-create;
        }
      }
      filter filter-name {
        enhanced-mode-override;
        fast-lookup-filter;
        instance-shared;
        interface-shared;
      interface-specific;
        term term-name {
          from {
            match-conditions;
          }
          then {
            action;
            action-modifiers;
          }
        only-at-create;
      filter filter-name {
      interface-specific;
        term term-name {
          from {
            match-conditions;
          }
          then {
            action;
            action-modifiers;
          }
        }
    policer policer-name {
      filter-specific;
      if-exceeding {
        (bandwidth-limit bps | bandwidth-percent percentage);
        burst-size-limit bytes;
      }
      logical-bandwidth-policer;
      logical-interface-policer;
      physical-interface-policer;
      then {
        policer-action;
      }
    }
    hierarchical-policer uid {
      aggregate {
        if-exceeding {
          bandwidth-limit-limit bps;
          burst-size-limit bytes;
```

 Copyright © 2018, Juniper Networks, Inc.

```
        }
        then {
          policer-action;
        }
      }
      premium {
        if-exceeding {
          bandwidth-limit bps;
          burst-size-limit bytes;
        }
        then {
          policer-action;
        }
      }
    }
    policer uid {
      filter-specific;
      if-exceeding {
        (bandwidth-limit bps | bandwidth-percent percentage);
        burst-size-limit bytes;
      }
      logical-bandwidth-policer;
      logical-interface-policer;
      physical-interface-policer;
      then {
        policer-action;
      }
    }
    three-color-policer uid {
      action {
        loss-priority high then discard;
      }
      logical-interface-policer;
      single-rate {
        (color-aware | color-blind);
        committed-burst-size bytes;
        committed-information-rate bps;
        excess-burst-size bytes;
      }
      two-rate {
        (color-aware | color-blind);
        committed-burst-size bytes;
        committed-information-rate bps;
        peak-burst-size bytes;
        peak-information-rate bps;
        }
      }
    }
  }
  interfaces interface-name {
    interface-set interface-set-name {
      interface interface-name {
        unit logical unit number {
          advisory-options {
            downstream-rate rate;
            upstream-rate rate;
```

```
            }
        }
    }
}
unit logical-unit-number {
    actual-transit-statistics;
    auto-configure {
        agent-circuit-identifier {
            dynamic-profile profile-name;
        }
        line-identity {
            include {
                accept-no-ids;
                circuit-id;
                remote-id;
            }
            dynamic-profile profile-name;
        }
    }
    encapsulation (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-cisco-nlpid |
        atm-tcc-vc-mux | atm-mlppp-llc | atm-nlpid | atm-ppp-llc | atm-ppp-vc-mux |
        atm-snap | atm-tcc-snap | atm-vc-mux | ether-over-atm-llc |
        ether-vpls-over-atm-llc | ether-vpls-over-fr | ether-vpls-over-ppp | ethernet |
        frame-relay-ccc | frame-relay-ppp | frame-relay-tcc | frame-relay-ether-type |
        frame-relay-ether-type-tcc | multilink-frame-relay-end-to-end | multilink-ppp |
        ppp-over-ether | ppp-over-ether-over-atm-llc | vlan-bridge | vlan-ccc | vlan-vci-ccc
        | vlan-tcc | vlan-vpls);
    family family {
        address address;
        filter {
            adf {
                counter;
                input-precedence precedence;
                not-mandatory;
                output-precedence precedence;
                rule rule-value;
            }
            input filter-name (
                precedence precedence;
                shared-name filter-shared-name;
            }
            output filter-name {
                precedence precedence;
                shared-name filter-shared-name;
            }
        }
        rpf-check {
            fail-filter filter-name;
            mode loose;
        }
        service {
            input {
                service-set service-set-name {
                    service-filter filter-name;
                }
                post-service-filter filter-name;
```

 Copyright © 2018, Juniper Networks, Inc.

```
        }
        input-vlan-map {
          inner-tag-protocol-id tpid;
          inner-vlan-id number;
          (push | swap);
          tag-protocol-id tpid;
          vlan-id number;
        }
        output {
          service-set service-set-name {
            service-filter filter-name;
          }
        }
        output-vlan-map {
          inner-tag-protocol-id tpid;
          inner-vlan-id number;
          (pop | swap);
          tag-protocol-id tpid;
          vlan-id number;
        }
        pcef  pcef-profile-name {
          activate rule-name | activate-all;
        }
      }
      unnumbered-address interface-name <preferred-source-address address>;
    }
    filter {
      input filter-name (
        shared-name filter-shared-name;
      }
      output filter-name {
        shared-name filter-shared-name;
      }
    }
    host-prefix-only;
    ppp-options {
      aaa-options aaa-options-name;
      authentication [ authentication-protocols ];
      chap {
        challenge-length minimum minimum-length maximum maximum-length;
        local-name name;
      }
      ignore-magic-number-mismatch;
      initiate-ncp (dual-stack-passive | ipv6 | ip)
      ipcp-suggest-dns-option;
      mru size;
      mtu (size | use-lower-layer);
      on-demand-ip-address;
      pap;
      peer-ip-address-optional;
      local-authentication {
        password password;
        username-include {
          circuit-id;
          delimiter character;
          domain-name name;
```

```
                mac-address;
                remote-id;
              }
            }
          }
          vlan-id number;
          vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
        }
      }
      interfaces {
        demux0 {...}
      }
      interfaces {
        pp0 {...}
      }
      policy-options {
        prefix-list uid {
          ip-addresses;
          dynamic-db;
        }
      }
      predefined-variable-defaults predefined-variable <variable-option> default-value;
      protocols {
        igmp {
          interface interface-name {
            accounting;
            disable;
            group-limit limit;
            group-policy;
            group-threshold value;
            immediate-leave
            log-interval seconds;
            no-accounting;
            oif-map;
            passive;
            promiscuous-mode;
            ssm-map ssm-map-name;
            ssm-map-policy ssm-map-policy-name
            static {
              group group {
                source source;
              }
            }
            version version;
          }
        }
        mld {
          interface interface-name {
            (accounting | no-accounting);
            disable;
            group-limit limit;
            group-policy;
            group-threshold value;
            immediate-leave;
            log-interval seconds;
            oif-map;
```

 Copyright © 2018, Juniper Networks, Inc.

```
          passive;
          ssm-map ssm-map-name;
          ssm-map-policy ssm-map-policy-name;
          static {
            group multicast-group-address {
              exclude;
              group-count number;
              group-increment increment;
              source ip-address {
                source-count number;
                source-increment increment;
              }
            }
          }
          version version;
        }
      }
      router-advertisement {
        interface interface-name {
          current-hop-limit number;
          default-lifetime seconds;
          (managed-configuration | no-managed-configuration);
          max-advertisement-interval seconds;
          min-advertisement-interval seconds;
          (other-stateful-configuration | no-other-stateful-configuration);
          prefix prefix;
          reachable-time milliseconds;
          retransmit-timer milliseconds;
        }
      }
    }
    routing-instances routing-instance-name {
      interface interface-name;
      routing-options {
        access {
          route prefix {
            next-hop next-hop;
            metric route-cost;
            preference route-distance;
            tag route-tag;
            tag2 route-tag2;
          }
        }
        access-internal {
          route subscriber-ip-address {
            qualified-next-hop underlying-interface {
              mac-address address;
            }
          }
        }
        multicast {
          interface interface-name {
            no-qos-adjust;
          }
        }
      }
```

```
    rib routing-table-name {
        access {
            route prefix {
                next-hop next-hop;
                metric route-cost;
                preference route-distance;
                tag route-tag;
                tag2 route-tag2;
            }
        }
        access-internal {
            route subscriber-ip-address {
                qualified-next-hop underlying-interface {
                    mac-address address;
                }
            }
        }
    }
}
routing-options {
    access {
        route prefix {
            next-hop next-hop;
            metric route-cost;
            preference route-distance;
            tag route-tag;
            tag2 route-tag2;
        }
    }
    access-internal {
        route subscriber-ip-address {
            qualified-next-hop underlying-interface {
                mac-address address;
            }
        }
    }
    multicast {
        interface interface-name {
            no-qos-adjust;
        }
    }
}
services {
    captive-portal-content-delivery {
        rule name {
            match-direction (input | input-output | output);
            term name {
                from {
                    applications application-name {
                        application-protocol type;
                        destination-port port-type;
                        protocol ip-protocol-type;
                        source-port port-type;
                    }
                    destination-address name <except>;
                    destination-address-range low minimum-value high maximum-value <except>;
```

 Copyright © 2018, Juniper Networks, Inc.

```
                destination-prefix-list name <except>;
              }
              then {
                accept;
                redirect url;
                rewrite destination-address address <destination-port port-number>;
                syslog;
              }
            }
          }
        }
      }
      variables {
        variable-name {
          default-value default-value;
          equals expression;
          mandatory;
          uid;
          uid-reference;
        }
      }
    }
  }
```

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 9.2.
Support at the **filter, policer, hierarchical-policer, three-color-policer, and policy options** hierarchy levels introduced in Junos OS Release 11.4.

Description Create dynamic profiles for use with DHCP or PPP client access.

Options ***profile-name***—Name of the dynamic profile; string of up to 80 alphanumeric characters.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

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

Related Documentation

- *Configuring a Basic Dynamic Profile*
- *Configuring Dynamic VLANs Based on Agent Circuit Identifier Information*
- *Dynamic Profiles Overview*

Copyright © 2018, Juniper Networks, Inc.

e1-options

Syntax

```
e1-options {
  bert-algorithm algorithm;
  bert-error-rate rate;
  bert-period seconds;
  fcs (16 | 32);
  framing (g704 | g704-no-crc4 | unframed);
  idle-cycle-flag (flags | ones);
  invert-data;
  loopback (local | remote);
  start-end-flag (filler | shared);
  timeslots time-slot-range;
}
```

Hierarchy Level

[edit interfaces *interface-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers.

Description

Configure E1-specific physical interface properties.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Channelized E1 IQ and IQE Interfaces Overview*
- *Channelized STM1 Interfaces Overview*
- *E1 Interfaces Overview*
- *T1 Interfaces Overview*

Copyright © 2018, Juniper Networks, Inc.

e3-options

Syntax

```
e3-options {
  atm-encapsulation (direct | plcp);
  bert-algorithm algorithm;
  bert-error-rate rate;
  bert-period seconds;
  buildout feet;
  compatibility-mode (digital-link | kentrox | larscom) <subrate value>;
  fcs (16 | 32);
  framing (g.751 | g.832);
  idle-cycle-flag value;
  invert-data;
  loopback (local | remote);
  (payload-scrambler | no-payload-scrambler);
  start-end-flag value;
  (unframed | no-unframed);
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure E3-specific physical interface properties.

For ATM1 interfaces, you can configure a subset of E3 options statements.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *E3 Interfaces Overview*
- *T3 Interfaces Overview*
- atm-options on page 444

east-interface

Syntax

```
east-interface {
  node-id mac-address;
  control-channel channel-name {
    vlan number;
    interface name interface-name
  }
  interface-none
  ring-protection-link-end;
}
```

Hierarchy Level

[edit protocols protection-group ethernet-ring *ring-name*]

Release Information

Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

Description

Define one of the two interface ports for Ethernet ring protection, the other being defined by the **west-interface** statement at the same hierarchy level. The interface must use the control channel's logical interface name. The control channel is a dedicated VLAN channel for the ring port.

EX Series switches do not use the node-id statement--the node ID is automatically configured on the switches using the MAC address.

NOTE: Always configure this port first, before configuring the **west-interface** statement.

NOTE: The Node ID is not configurable on EX Series switches. The node ID is automatically configured using the MAC address.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Ethernet Ring Protection Switching Overview*
- *Ethernet Ring Protection Using Ring Instances for Load Balancing*
- west-interface on page 1105
- ethernet-ring on page 595

 Copyright © 2018, Juniper Networks, Inc.

- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

egress-policer-overhead

Syntax egress-policer-overhead *bytes*;

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number*]

Release Information Statement introduced before Junos OS Release 11.1.

Description Add the specified number of bytes to the actual length of an Ethernet frame when determining the actions of Layer 2 policers, MAC policers, or queue rate limits applied to output traffic on the line card. You can configure egress policer overhead to account for egress *shaping* overhead bytes added to output traffic on the line card.

On M Series and T Series routers, this statement is supported on Gigabit Ethernet Intelligent Queuing 2 (IQ2) PICs and Enhanced IQ2 (IQ2E) PICs. On MX Series routers, this statement is supported for interfaces configured on Dense Port Concentrators (DPCs).

NOTE: **This statement is not supported on Modular Interface Cards (MICs) or Modular Port Concentrators (MPCs) in MX Series routers.**

Options *bytes*—Number of bytes added to a packet exiting an interface.

Range: 0–255 bytes

Default: 0

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

Related Documentation

- *egress-shaping-overhead*
- *Policer Overhead to Account for Rate Shaping Overview*
- *Example: Configuring Policer Overhead to Account for Rate Shaping*
- *Configuring a Policer Overhead*
- *CoS on Enhanced IQ2 PICs Overview*

Copyright © 2018, Juniper Networks, Inc.

encapsulation (Container Interface)

Syntax encapsulation (cisco-hdlc | ppp);

Hierarchy Level [edit interfaces ci*n*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Container link-layer encapsulation type.

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

ppp—Use serial PPP encapsulation.

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

Related Documentation

- *Displaying APS Using a Container Interface with ATM Encapsulation*
- *Configuring Container Interfaces for APS on SONET Links*

 Copyright © 2018, Juniper Networks, Inc.

encapsulation (Logical Interface)

Syntax

encapsulation (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-cisco-nlpid | atm-mlppp-llc | atm-nlpid | atm-ppp-llc | atm-ppp-vc-mux | atm-snap | atm-tcc-snap | atm-tcc-vc-mux | atm-vc-mux | ether-over-atm-llc | ether-vpls-over-atm-llc | ether-vpls-over-fr | ether-vpls-over-ppp | ethernet | ethernet-ccc | ethernet-vpls | ethernet-vpls-fr | frame-relay-ccc | frame-relay-ether-type | frame-relay-ether-type-tcc | frame-relay-ppp | frame-relay-tcc | gre-fragmentation | multilink-frame-relay-end-to-end | multilink-ppp | ppp-over-ether | ppp-over-ether-over-atm-llc | vlan-bridge | vlan-ccc | vlan-vci-ccc | vlan-tcc | vlan-vpls | vxlan);

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers (**ethernet**,**vlan-ccc**, and **vlan-tcc** options only).
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers. Only the **atm-ccc-cell-relay** and **atm-ccc-vc-mux** options are supported on ACX Series routers.
Statement introduced in Junos OS Release 17.3R1 for QFX10000 Series switches (**ethernet-ccc** and **vlan-ccc** options only).

Description

Configure a logical link-layer encapsulation type. Not all encapsulation types are supported on the switches. See the switch CLI.

Options

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

atm-ccc-vc-mux—Use ATM virtual circuit (VC) multiplex encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

atm-cisco-nlpid—Use Cisco ATM network layer protocol identifier (NLPID) encapsulation. When you use this encapsulation type, you can configure the **inet** family only.

atm-mlppp-llc—For ATM2 IQ interfaces only, use Multilink Point-to-Point (MLPPP) over AAL5 LLC. For this encapsulation type, your router must be equipped with a Link Services or Voice Services PIC. MLPPP over ATM encapsulation is not supported on ATM2 IQ OC48 interfaces.

atm-nlpid—Use ATM NLPID encapsulation. When you use this encapsulation type, you can configure the **inet** family only.

atm-ppp-llc—(ATM2 IQ interfaces and MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP only) Use PPP over AAL5 LLC encapsulation.

atm-ppp-vc-mux—(ATM2 IQ interfaces and MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP only) Use PPP over ATM AAL5 multiplex encapsulation.

atm-snap—(All interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) Use ATM subnetwork attachment point (SNAP) encapsulation.

atm-tcc-snap—Use ATM SNAP encapsulation on translational cross-connect (TCC) circuits.

atm-tcc-vc-mux—Use ATM VC multiplex encapsulation on TCC circuits. When you use this encapsulation type, you can configure the **tcc** family only.

atm-vc-mux—(All interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) Use ATM VC multiplex encapsulation. When you use this encapsulation type, you can configure the **inet** family only.

ether-over-atm-llc—(All IP interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) For interfaces that carry IP traffic, use Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces.

ether-vpls-over-atm-llc—For ATM2 IQ interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.

ether-vpls-over-fr—For E1, T1, E3, T3, and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over Frame Relay encapsulation to support Bridged Ethernet over Frame Relay encapsulated TDM interfaces for VPLS applications, per RFC 2427, *Multiprotocol Interconnect over Frame Relay*.

NOTE: **The SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP, the Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP, and the DS3/E3 MIC do not support Ethernet over Frame Relay encapsulation.**

ether-vpls-over-ppp—For E1, T1, E3, T3, and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over Point-to-Point Protocol (PPP) encapsulation to support Bridged Ethernet over PPP-encapsulated TDM interfaces for VPLS applications.

ethernet—Use Ethernet II encapsulation (as described in RFC 894, *A Standard for the Transmission of IP Datagrams over Ethernet Networks*).

ethernet-ccc—Use Ethernet CCC encapsulation on Ethernet interfaces.

ethernet-vpls—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard Tag Protocol ID (TPID) values.

 Copyright © 2018, Juniper Networks, Inc.

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

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

frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

frame-relay-ether-type—Use Frame Relay ether type encapsulation for compatibility with Cisco Frame Relay. The physical interface must be configured with flexible-frame-relay encapsulation.

frame-relay-ether-type-tcc—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect different media. The physical interface must be configured with flexible-frame-relay encapsulation.

frame-relay-ppp—Use PPP over Frame Relay circuits. When you use this encapsulation type, you can configure the **ppp** family only.

frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits for connecting different media. When you use this encapsulation type, you can configure the **tcc** family only.

gre-fragmentation—For adaptive services interfaces only, use GRE fragmentation encapsulation to enable fragmentation of IPv4 packets in GRE tunnels. This encapsulation clears the do not fragment (DF) bit in the packet header. If the packet's size exceeds the tunnel's maximum transmission unit (MTU) value, the packet is fragmented before encapsulation.

multilink-frame-relay-end-to-end—Use MLFR FRF.15 encapsulation. This encapsulation is used only on multilink, link services, and voice services interfaces and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.

multilink-ppp—Use MLPPP encapsulation. This encapsulation is used only on multilink, link services, and voice services interfaces and their constituent T1 or E1 interfaces.

ppp-over-ether—Use PPP over Ethernet encapsulation to configure an underlying Ethernet interface for a dynamic PPPoE logical interface on M120 and M320 routers with Intelligent Queuing 2 (IQ2) PICs, and on MX Series routers with MPCs.

ppp-over-ether-over-atm-llc—(MX Series routers with MPCs using the ATM MIC with SFP only) For underlying ATM interfaces, use PPP over Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure the interface address. Instead, configure the interface address on the PPP interface.

Copyright © 2018, Juniper Networks, Inc.

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

vlan-ccc—Use Ethernet virtual LAN (VLAN) encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

vlan-vci-ccc—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

vlan-tcc—Use Ethernet VLAN encapsulation on TCC circuits. When you use this encapsulation type, you can configure the **tcc** family only.

vlan-vpls—Use Ethernet VLAN encapsulation on VPLS circuits.

vxlan—Use VXLAN data plane encapsulation for EVPN.

Required Privilege Level

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

Related Documentation

- *Configuring Layer 2 Switching Cross-Connects Using CCC*
- *Configuring the Encapsulation for Layer 2 Switching TCCs*
- Configuring Interface Encapsulation on Logical Interfaces on page 190
- *Configuring the CCC Encapsulation for LSP Tunnel Cross-Connects*
- Circuit and Translational Cross-Connects Overview on page 275
- *Identifying the Access Concentrator*
- *Configuring ATM Interface Encapsulation*
- *Configuring VLAN and Extended VLAN Encapsulation*
- Configuring ATM-to-Ethernet Interworking on page 282
- Configuring Interface Encapsulation on PTX Series Packet Transport Routers on page 132
- *Configuring CCC Encapsulation for Layer 2 VPNs*
- *Configuring TCC Encapsulation for Layer 2 VPNs and Layer 2 Circuits*
- *Configuring ATM for Subscriber Access*
- *Understanding CoS on ATM IMA Pseudowire Interfaces Overview*
- *Configuring Policing on an ATM IMA Pseudowire*

 Copyright © 2018, Juniper Networks, Inc.

encapsulation

List of Syntax

Syntax for Physical Interfaces: M Series, MX Series, QFX Series, T Series, PTX Series on page 577
Syntax for Logical Interfaces: SRX Series on page 577

Syntax for Physical Interfaces: M Series, MX Series, QFX Series, T Series, PTX Series

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

Syntax for Logical Interfaces: SRX Series

```
encapsulation (ether-vpls-ppp | ethernet-bridge | ethernet-ccc | ethernet-tcc | ethernet-vpls
  | extended-frame-relay-ccc | extended-frame-relay-tcc | extended-vlan-bridge |
  extended-vlan-ccc | extended-vlan-tcc | extended-vlan-vpls | frame-relay-port-ccc |
  vlan-ccc | vlan-vpls);
```

Physical Interfaces: M Series, MX Series, QFX Series, T Series, PTX Series

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

Logical Interfaces: SRX Series

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

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.5.
Statement introduced in Junos OS Release 11.1 for EX Series switches.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers (**flexible-ethernet-services**, **ethernet-ccc**, and **ethernet-tcc** options only).

Description

For M Series, MX Series, QFX Series, T Series, PTX Series, specify the physical link-layer encapsulation type.

For SRX Series, specify logical link layer encapsulation.

NOTE: Not all encapsulation types are supported on the switches. See the switch CLI.

Default

ppp—Use serial PPP encapsulation.

Copyright © 2018, Juniper Networks, Inc.

Physical Interface Options and Logical Interface Options

[Warning: element unresolved in stylesheets: <title> (in <config-options>). This is probably a new element that is not yet supported in the stylesheets.]

Physical Interface Options and Logical Interface Options

For physical interfaces:

NOTE: Frame Relay, ATM, PPP, SONET, and SATSOP options are not supported on EX Series switches.

- **atm-ccc-cell-relay**—Use ATM cell-relay encapsulation.
- **atm-pvc**—Defined in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*. When you configure physical ATM interfaces with ATM PVC encapsulation, an RFC 2684-compliant ATM Adaptation Layer 5 (AAL5) tunnel is set up to route the ATM cells over a Multiprotocol Label Switching (MPLS) path that is typically established between two MPLS-capable routers using the Label Distribution Protocol (LDP).
- **cisco-hdlc**—Use Cisco-compatible High-Level Data Link Control (HDLC) framing. E1, E3, SONET/SDH, T1, and T3 interfaces can use Cisco HDLC encapsulation. Two related versions are supported:
  - CCC version (**cisco-hdlc-ccc**)—The logical interface does not require an encapsulation statement. When you use this encapsulation type, you can configure the **ccc** family only.
  - TCC version (**cisco-hdlc-tcc**)—Similar to CCC and has the same configuration restrictions, but used for circuits with different media on either side of the connection.
- **cisco-hdlc-ccc**—Use Cisco-compatible HDLC framing on CCC circuits.
- **cisco-hdlc-tcc**—Use Cisco-compatible HDLC framing on TCC circuits for connecting different media.
- **ethernet-bridge**—Use Ethernet bridge encapsulation on Ethernet interfaces that have bridging enabled and that must accept all packets.
- **ethernet-over-atm**—For interfaces that carry IPv4 traffic, use Ethernet over ATM encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces. As defined in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*, this encapsulation type allows ATM interfaces to connect to devices that support only bridge protocol data units (BPDUs). Junos OS does not completely support bridging, but accepts BPDU packets as a default gateway. If you use the router as an edge device, then the router acts as a default gateway. It accepts Ethernet LLC/SNAP frames with IP or ARP in the payload, and drops the rest. For packets destined to the Ethernet LAN, a route lookup is done using the destination IP address. If the route lookup yields a full address match, the packet is encapsulated with an LLC/SNAP and MAC header, and the packet is forwarded to the ATM interface.
- **ethernet-tcc**—For interfaces that carry IPv4 traffic, use Ethernet TCC encapsulation on interfaces that must accept packets carrying standard TPID values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, TCC is not supported.

 Copyright © 2018, Juniper Networks, Inc.

- **ethernet-vpls**—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard TPID values. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.
- **ethernet-vpls-fr**—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer Layer 2 Frame Relay connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use the MAC address to forward the packet into a given VPLS instance.
- **ethernet-vpls-ppp**—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer Layer 2 PPP connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use it to forward the packet into a given VPLS instance.
- **ether-vpls-over-atm-llc**—For ATM intelligent queuing (IQ) interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.
- **extended-frame-relay-ccc**—Use Frame Relay encapsulation on CCC circuits. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to CCC. When you use this encapsulation type, you can configure the **ccc** family only.
- **extended-frame-relay-ether-type-tcc**—Use extended Frame Relay ether type TCC for Cisco-compatible Frame Relay for DLCIs 1 through 1022. This encapsulation type is used for circuits with different media on either side of the connection.
- **extended-frame-relay-tcc**—Use Frame Relay encapsulation on TCC circuits to connect different media. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to TCC.
- **extended-vlan-bridge**—Use extended VLAN bridge encapsulation on Ethernet interfaces that have IEEE 802.1Q VLAN tagging and bridging enabled and that must accept packets carrying TPID 0x8100 or a user-defined TPID.
- **extended-vlan-ccc**—Use extended VLAN encapsulation on CCC circuits with Gigabit Ethernet and 4-port Fast Ethernet interfaces that must accept packets carrying 802.1Q values. Extended VLAN CCC encapsulation supports TPIDs 0x8100, 0x9100, and 0x9901. When you use this encapsulation type, you can configure the **ccc** family only. For 8-port, 12-port, and 48-port Fast Ethernet PICs, extended VLAN CCC is not supported. For 4-port Gigabit Ethernet PICs, extended VLAN CCC is not supported.
- **extended-vlan-tcc**—For interfaces that carry IPv4 traffic, use extended VLAN encapsulation on TCC circuits with Gigabit Ethernet interfaces on which you want to use 802.1Q tagging. For 4-port Gigabit Ethernet PICs, extended VLAN TCC is not supported.

Copyright © 2018, Juniper Networks, Inc.

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

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

- **flexible-ethernet-services**—For Gigabit Ethernet IQ interfaces and Gigabit Ethernet PICs with small form-factor pluggable transceivers (SFPs) (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), and for Gigabit Ethernet interfaces, use flexible Ethernet services encapsulation when you want to configure multiple per-unit Ethernet encapsulations. Aggregated Ethernet bundles can use this encapsulation type. This encapsulation type allows you to configure any combination of route, TCC, CCC, Layer 2 virtual private networks (VPNs), and VPLS encapsulations on a single physical port. If you configure flexible Ethernet services encapsulation on the physical interface, VLAN IDs from 1 through 511 are no longer reserved for normal VLANs.
- **flexible-frame-relay**—For IQ interfaces only, use flexible Frame Relay encapsulation when you want to configure multiple per-unit Frame Relay encapsulations. This encapsulation type allows you to configure any combination of TCC, CCC, and standard Frame Relay encapsulations on a single physical port. Also, each logical interface can have any DLCI value from 1 through 1022.
- **frame-relay**—Use Frame Relay encapsulation is defined in RFC 1490, *Multiprotocol Interconnect over Frame Relay*. E1, E3, link services, SONET/SDH, T1, T3, and voice services interfaces can use Frame Relay encapsulation.
- **frame-relay-ccc**—Use Frame Relay encapsulation on CCC circuits. This encapsulation is same as standard Frame Relay for DLCIs 0 through 511. DLCIs 512 through 1022 are dedicated to CCC. The logical interface must also have **frame-relay-ccc** encapsulation. When you use this encapsulation type, you can configure the **ccc** family only.
- **frame-relay-ether-type**—Use Frame Relay ether type encapsulation for compatibility with the Cisco Frame Relay. IETF frame relay encapsulation identifies the payload format using NLPID and SNAP formats. Cisco-compatible Frame Relay encapsulation uses the Ethernet type to identify the type of payload.

NOTE: **When the encapsulation type is set to Cisco-compatible Frame Relay encapsulation, ensure that the LMI type is set to ANSI or Q933-A.**

- **frame-relay-ether-type-tcc**—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect different media. This encapsulation is Cisco-compatible Frame Relay for DLCIs 0 through 511. DLCIs 512 through 1022 are dedicated to TCC.

 Copyright © 2018, Juniper Networks, Inc.

- **frame-relay-port-ccc**—Use Frame Relay port CCC encapsulation to transparently carry all the DLCIs between two customer edge (CE) routers without explicitly configuring each DLCI on the two provider edge (PE) routers with Frame Relay transport. The connection between the two CE routers can be either user-to-network interface (UNI) or network-to-network interface (NNI); this is completely transparent to the PE routers. When you use this encapsulation type, you can configure the **ccc** family only.
- **frame-relay-tcc**—This encapsulation is similar to Frame Relay CCC and has the same configuration restrictions, but used for circuits with different media on either side of the connection.
- **generic-services**—Use generic services encapsulation for services with a hierarchical scheduler.
- **multilink-frame-relay-uni-nni**—Use MLFR UNI NNI encapsulation. This encapsulation is used on link services, voice services interfaces functioning as FRF.16 bundles, and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.
-
- **ppp**—Use serial PPP encapsulation. This encapsulation is defined in RFC 1661, *The Point-to-Point Protocol (PPP) for the Transmission of Multiprotocol Datagrams over Point-to-Point Links*. PPP is the default encapsulation type for physical interfaces. E1, E3, SONET/SDH, T1, and T3 interfaces can use PPP encapsulation.
- **ppp-ccc**—Use serial PPP encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.
- **ppp-tcc**—Use serial PPP encapsulation on TCC circuits for connecting different media. When you use this encapsulation type, you can configure the **tcc** family only.
- **vlan-ccc**—Use Ethernet VLAN encapsulation on CCC circuits. VLAN CCC encapsulation supports TPID 0x8100 only. When you use this encapsulation type, you can configure the **ccc** family only.

- **vlan-vci-ccc**—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only. All logical interfaces configured on the Ethernet interface must also have the encapsulation type set to **vlan-vci-ccc**.
- **vlan-vpls**—Use VLAN VPLS encapsulation on Ethernet interfaces with VLAN tagging and VPLS enabled. Interfaces with VLAN VPLS encapsulation accept packets carrying standard TPID values only. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.

NOTE:
- Label-switched interfaces (LSIs) do not support VLAN VPLS encapsulation. Therefore, you can only use VLAN VPLS encapsulation on a PE-router-to-CE-router interface and not a core-facing interface.
- Starting with Junos OS release 13.3, a commit error occurs when you configure **vlan-vpls** encapsulation on a physical interface and configure **family inet** on one of the logical units. Previously, it was possible to commit this invalid configuration.

For logical interfaces:

- **frame-relay**—Configure a Frame Relay encapsulation when the physical interface has multiple logical units, and the units are either point to point or multipoint.
- **multilink-frame-relay-uni-nni**—Link services interfaces functioning as FRF.16 bundles can use Multilink Frame Relay UNI NNI encapsulation.
- **ppp**—For normal mode (when the device is using only one ISDN B-channel per call). Point-to-Point Protocol is for communication between two computers using a serial interface.
- **ppp-over-ether**—This encapsulation is used for underlying interfaces of pp0 interfaces.

Required Privilege Level

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

 Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- *Understanding Physical Encapsulation on an Interface*
- Configuring Interface Encapsulation on Physical Interfaces on page 129
- *Configuring CCC Encapsulation for Layer 2 VPNs*
- *Configuring Layer 2 Switching Cross-Connects Using CCC*
- *Configuring TCC Encapsulation for Layer 2 VPNs and Layer 2 Circuits*
- *Configuring ATM Interface Encapsulation*
- Configuring ATM-to-Ethernet Interworking on page 282
- *Configuring VLAN and Extended VLAN Encapsulation*
- *Configuring VLAN and Extended VLAN Encapsulation*
- *Configuring Encapsulation for Layer 2 Wholesale VLAN Interfaces*
- *Configuring Interfaces for Layer 2 Circuits*
- Configuring Interface Encapsulation on PTX Series Packet Transport Routers on page 132
- *Configuring MPLS LSP Tunnel Cross-Connects Using CCC*
- *Configuring TCC*
- *Configuring VPLS Interface Encapsulation*
- *Configuring Interfaces for VPLS Routing*
- Defining the Encapsulation for Switching Cross-Connects on page 277
- *Configuring an MPLS-Based Layer 2 VPN (CLI Procedure)*

Copyright © 2018, Juniper Networks, Inc.

encoding

Syntax encoding (nrz | nrzi);

Hierarchy Level [edit interfaces *interface-name* serial-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For serial interfaces, set the line encoding format.

Default The default line encoding is non-return to zero (NRZ).

Options **nrz**—Use NRZ line encoding.

nrzi—Use non-return to zero inverted (NRZI) line encoding.

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

Related Documentation

- Configuring Serial Line Encoding on page 350

 Copyright © 2018, Juniper Networks, Inc.

epd-threshold (Logical Interface)

Syntax

epd-threshold *cells* plp1 *cells*;

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number*],
[edit interfaces *interface-name* unit *logical-unit-number* address *address* family *family* multipoint-destination *address*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* address *address* family *family* multipoint-destination *address*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description

For ATM2 IQ interfaces only, define the early packet discard (EPD) threshold on a VC. The EPD threshold is a limit on the number of transmit packets that can be queued. Packets that exceed the limit are discarded. For interfaces configured in trunk mode, you can also configure dual EPD thresholds depending on the packet loss priorities (PLPs).

Default

Approximately 1 percent of the available cell buffers. If shaping is enabled, the default EPD threshold is proportional to the shaping rate according to the following formula:

default epd-threshold = number of buffers * shaping rate / line rate

The minimum EPD threshold value is 48 cells. If the default EPD threshold formula results in an EPD threshold of less than 48 cells, the result will be ignored, and the minimum value of 48 cells will be used.

Options

cells—Maximum number of cells.

Range: For 1-port and 2-port OC12 interfaces, 48 through 425,984 cells

Required Privilege Level

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

Related Documentation

- *Configuring the ATM2 IQ EPD Threshold*
- *Configuring Two EPD Thresholds per Queue*

epd-threshold (Physical Interface)

Syntax epd-threshold *cells* plp1 *cells*;

Hierarchy Level [edit interfaces at-*fpc/pic/port* atm-options scheduler-maps *map-name* forwarding-class *class-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For ATM2 IQ interfaces only, define the EPD threshold on a VC. The EPD threshold is a limit on the number of transmit packets that can be queued. Packets that exceed the limit are discarded.

Default If you do not include either the **epd-threshold** or the **linear-red-profile** statement in the forwarding class configuration, the Junos OS uses an EPD threshold based on the available bandwidth and other parameters.

Options *cells*—Maximum number of cells.

Range: For 1-port and 2-port OC12 interfaces, 48 through 425,984 cells. For 1-port OC48 interfaces, 48 through 425,984 cells. For 2-port OC3, DS3, and E3 interfaces, 48 through 212,992 cells. For 4-port DS3 and E3 interfaces, 48 through 106,496 cells.

The **plp1** statement is explained separately.

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

Related Documentation

- *Configuring an ATM Scheduler Map*
- linear-red-profile on page 725

 Copyright © 2018, Juniper Networks, Inc.

es-options

Syntax

```
es-options {
  backup-interface interface-name;
}
```

Hierarchy Level [edit interfaces es-*fpc/pic/port*]

Release Information Statement introduced before Junos OS Release 7.4.

Description On ES interfaces, configure ES interface-specific interface properties.

The **backup-interface** statement is explained separately.

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

ethernet (Protocols OAM)

List of Syntax

Syntax: MX, T, ACX Series Routers, SRX Firewalls, M320 and EX Series Switches on page 588
Syntax: EX Series Switches and NFX Series Devices on page 591

Syntax: MX, T, ACX Series Routers, SRX Firewalls, M320 and EX Series Switches

```
ethernet {
  connectivity-fault-management {
    action-profile profile-name {
      default-actions {
        interface-down;
      }
    }
    performance-monitoring {
      delegate-server-processing;
      hardware-assisted-timestamping;
      hardware-assisted-keepalives;
      sla-iterator-profiles {
        profile-name {
          avg-fd-twoway-threshold;
          avg-ifdv-twoway-threshold;
          avg-flr-forward-threshold;
          avg-flr-backward-threshold;
          disable;
          calculation-weight {
            delay delay-weight;
            delay-variation delay-variation-weight;
          }
          cycle-time milliseconds;
          iteration-period connections;
          measurement-type (loss | statistical-frame-loss | two-way-delay);
        }
      }
    }
    linktrace {
      age (30m | 10m | 1m | 30s | 10s);
      path-database-size path-database-size;
    }
    maintenance-domain domain-name {
      level number;
      name-format (character-string | none | dns | mac+2octet);
      maintenance-association ma-name {
        short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
        protect-maintenance-association protect-ma-name;
        remote-maintenance-association remote-ma-name;
        continuity-check {
          convey-loss-threshold;
          hold-interval minutes;
          interface-status-tlv;
          interval (10m | 10s | 1m | 1s| 100ms);
          loss-threshold number;
          port-status-tlv;
        }
        mep mep-id {
```

 Copyright © 2018, Juniper Networks, Inc.

```
                auto-discovery;
                direction (up | down);
                interface interface-name (protect | working);
                lowest-priority-defect (all-defects | err-xcon | mac-rem-err-xcon | no-defect |
                  rem-err-xcon | xcon );
                priority number;
                remote-mep mep-id {
                  action-profile profile-name;
                  sla-iterator-profile profile-name {
                    data-tlv-size size;
                    iteration-count count-value;
                    priority priority-value;
                  }
                }
              }
            }
          }
        }
        evcs evc-id {
          evc-protocol cfm management-domain domain-id (management-association
            association-id | vpls (routing-instance instance-id);
          remote-uni-count count;
          multipoint-to-multipoint;
        }
        link-fault-management {
          action-profile profile-name {
            action {
              link-down;
              send-critical-event;
              syslog;
            }
            event {
              link-adjacency-loss;
              link-event-rate {
                frame-error count;
                frame-period count;
                frame-period-summary count;
                symbol-period count;
              }
              protocol-down;
            }
          }
          interface interface-name {
            apply-action-profile;
            link-discovery (active | passive);
            loopback-tracking;
            pdu-interval interval;
            pdu-threshold threshold-value;
            remote-loopback;
            event-thresholds {
              frame-error count;
              frame-period count;
              frame-period-summary count;
              symbol-period count;
            }
            negotiation-options {
```

Copyright © 2018, Juniper Networks, Inc.

```
            allow-remote-loopback;
            no-allow-link-events;
          }
        }
      }
      lmi {
        status-counter count;
        polling-verification-timer value;
        interface name {
          uni-id uni-name;
          status-counter number;
          polling-verification-timer value;
          evc-map-type (all-to-one-bundling | bundling | service-multiplexing);
          evc evc-name {
            default-evc;
            vlan-list vlan-id-list;
          }
        }
      }
    }
```


Copyright © 2018, Juniper Networks, Inc.

Syntax: EX Series Switches and NFX Series Devices

```
ethernet {
  connectivity-fault-management {
    action-profile profile-name {
      action {
        interface-down;
      }
      default-actions {
        interface-down;
      }
      event {
        adjacency-loss;
      }
    }
    esp-traceoptions {
      file filename <files number> <no-stamp> <replace> <size size> <world-readable |
        no-world-readable>;
      flag (all |error | esp | interface | krt | lib |normal |task |timer);
    }
    linktrace {
      age (30m | 10m | 1m | 30s | 10s);
      path-database-size path-database-size;
    }
    maintenance-domain domain-name {
      level number;
      mip-half-function (none | default |explicit);
      name-format (character-string | none | dns | mac+2oct);
      maintenance-association ma-name {
        continuity-check {
          hold-interval minutes;
          interface-status-tlv;
          interval (10m | 10s | 1m | 1s| 100ms);
          loss-threshold number;
          port-status-tlv;
        }
        mep mep-id {
          auto-discovery;
          direction down;
          interface interface-name;
          priority
          remote-mep mep-id {
            action-profile profile-name;
            sla-iterator-profile profile-name {
              data-tlv-size size;
              iteration-count count-value;
              priority priority-value;
            }
          }
        }
        short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
      }
    }
    performance-monitoring {
      sla-iterator-profiles {
        profile-name {
          calculation-weight {
            delay delay-value;
```

```
                        delay-variation delay-variation-value;
                      }
                      cycle-time cycle-time-value;
                      iteration-period iteration-period-value;
                      measurement-type two-way-delay;
                      passive;
                    }
                  }
                }
                traceoptions {
                  file filename <files number> <match regex> <size size> <world-readable |
                    no-world-readable>;
                  flag flag ;
                  no-remote-trace;
                }
              }
              link-fault-management {
                action-profile profile-name;
                  action {
                    syslog;
                    link-down;
                  }
                  event {
                    link-adjacency-loss;
                    link-event-rate {
                      frame-error  count;
                      frame-period count;
                      frame-period-summary  count;
                      symbol-period count;
                    }
                  }
                interface interface-name {
                  link-discovery (active | passive);
                  pdu-interval interval;
                  pdu-threshold threshold-value;
                  remote-loopback;
                  event-thresholds  {
                    frame-error count;
                    frame-period count;
                    frame-period-summary  count;
                    symbol-period count;
                  }
                  negotiation-options {
                    allow-remote-loopback;
                    no-allow-link-events;
                  }
                }
                traceoptions {
                  file filename <files number> <match regex> <size size> <world-readable |
                    no-world-readable>;
                  flag flag ;
                  no-remote-trace;
                }
              }
            }
```

 Copyright © 2018, Juniper Networks, Inc.

Hierarchy Level [edit protocols oam]

Release Information Statement introduced in Junos OS Release 8.2 for MX, T, ACX Series routers, SRX firewalls, M320 and EX Series switches.
Statement introduced in Junos OS Release 9.4 for EX Series switches and NFX Series devices.
connectivity-fault-management introduced in Junos OS Release 10.2 for EX Series switches.

Description Provide IEEE 802.3ah Operation, Administration, and Maintenance (OAM) support for Ethernet interfaces or configure connectivity fault management (CFM) for IEEE 802.1ag Operation, Administration, and Management (OAM) support.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *Enabling IEEE 802.3ah OAM Support*
- *Example: Configuring Ethernet OAM Link Fault Management*

ethernet-policer-profile

Syntax

```
ethernet-policer-profile {
  input-priority-map {
    ieee802.1p premium [ values ];
  }
  output-priority-map {
    classifier {
      premium {
        forwarding-class class-name {
          loss-priority (high | low);
        }
      }
    }
  }
  policer cos-policer-name {
    aggregate {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
    premium {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
  }
}
```

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile],
[edit interfaces *interface-name* aggregated-ether-options ethernet-switch-profile]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

NOTE: On QFX Series standalone switches, this statement hierarchy is only supported on the Enhanced Layer 2 Switching CLI.

For Gigabit Ethernet IQ, 10-Gigabit Ethernet, Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), and 100-Gigabit Ethernet Type 5 PIC with CFP, configure a class of service (CoS)-based policer. Policing applies to the inner VLAN identifiers, not to the outer tag. For Gigabit Ethernet interfaces with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), the **premium** policer is not supported.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- *Configuring Gigabit Ethernet Policers*

ethernet-ring

Syntax

```
ethernet-ring ring-name {
  control-vlan (vlan-id | vlan-name);
  data-channel {
    vlan number
  }
  east-interface {
    control-channel channel-name {
      vlan number;
      interface name interface-name
    }
  }
  guard-interval number;
  node-id mac-address;
  restore-interval number;
  ring-protection-link-owner;
  west-interface {
    control-channel channel-name {
      vlan number;
    }
  }
}
```

Hierarchy Level

[edit protocols protection-group]

Release Information

Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

Description

For Ethernet PICs on MX Series routers or for EX Series switches, , specify the Ethernet ring in an Ethernet ring protection switching configuration.

Options

ring-name—Name of the Ethernet protection ring.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Ethernet Ring Protection Switching Overview*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

Copyright © 2018, Juniper Networks, Inc.

ethernet-switch-profile

Syntax

```
ethernet-switch-profile {
  ethernet-policer-profile {
    input-priority-map {
      ieee802.1p premium [values];
    }
    output-priority-map {
      classifier {
        premium {
          forwarding-class class-name {
            loss-priority (high | low);
          }
        }
      }
    }
    policer cos-policer-name {
      aggregate {
        bandwidth-limit bps;
        burst-size-limit bytes;
      }
      premium {
        bandwidth-limit bps;
        burst-size-limit bytes;
      }
    }
    storm-control storm-control-profile;
    tag-protocol-id tpid;
  }
  mac-learn-enable;
}
```

Hierarchy Level

[edit interfaces *interface-name* gigether-options],
[edit interfaces *interface-name* aggregated-ether-options],
[edit interfaces *interface-name* aggregated-ether-options],
[edit interfaces *interface-name* ether-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.
Statement introduced in Junos OS Release 13.2 for the QFX Series.
Statement introduced in Junos OS Release 13.2X50-D15 for the EX Series switches.

Description

NOTE: On QFX Series standalone switches, the **ethernet-policer-profile** CLI hierarchy and the **mac-learn-enable** statement are supported only on the Enhanced Layer 2 Switching CLI.

For Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2 and IQ2-E, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC, aggregated Ethernet with Gigabit Ethernet IQ interfaces, the built-in Gigabit Ethernet port on the M7i router); 100-Gigabit

Copyright © 2018, Juniper Networks, Inc.

Ethernet Type 5 PIC with CFP; and Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces on EX Series switches, configure VLAN tag and MAC address accounting and filtering properties.

The remaining statements are explained separately. See CLI Explorer.

NOTE: When you gather interfaces into a bridge domain, the **no-mac-learn-enable** statement at the **[edit interfaces *interface-name* gigether-options ethernet-switch-profile]** hierarchy level is not supported. You must use the **no-mac-learning** statement at the **[edit bridge-domains *bridge-domain-name* bridge-options interface *interface-name*]** hierarchy level to disable MAC learning on an interface in a bridge domain. For information on disabling MAC learning for a bridge domain, see the *MX Series Layer 2 Configuration Guide*.

Default If the **ethernet-switch-profile** statement is not configured, Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router) behave like Gigabit Ethernet interfaces.

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

Related Documentation

- *Configuring Gigabit Ethernet Policers*
- *Configuring MAC Address Filtering*
- *Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview*
- *Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support (CLI Procedure)*

Copyright © 2018, Juniper Networks, Inc.

eui-64

Syntax eui-64;

Hierarchy Level [edit interfaces *interface-name* unit *number* family inet6 address *address*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description For interfaces that carry IP version 6 (IPv6) traffic, automatically generate the host number portion of interface addresses.

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

Related Documentation

- Configuring the Interface Address on page 207

Copyright © 2018, Juniper Networks, Inc.

evcs

Syntax

```
evcs evc-id {
  evc-protocol cfm;
  remote-uni-count count;
  multipoint-to-multipoint;
}
```

Hierarchy Level [edit protocols oam ethernet]

Release Information Statement introduced in Junos OS Release 9.5.

Description On MX Series routers with **ge**, **xe**, or **ae** interfaces, configure an OAM Ethernet virtual connection.

Options **remote-uni-count** ***count***—(Optional) Specify the number of remote UNIs in the EVC configuration, the default is 1.

multipoint-to-multipoint—(Optional) Specify multiple points in the EVC configuration, the default is point-to-point if **remote-uni-count** is 1.

Remaining options are explained separately.

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

Related Documentation

- *Configuring Ethernet Local Management Interface*
- lmi (Ethernet OAM) on page 740

event (LFM)

List of Syntax

Syntax: MX, M, T, ACX Series Routers, SRX Firewalls and EX Series Switches on page 600
Syntax: EX Series Switches and NFX Series Devices on page 600

Syntax: MX, M, T, ACX Series Routers, SRX Firewalls and EX Series Switches

```
event {
  link-adjacency-loss;
  link-event-rate {
    frame-error count;
    frame-period count;
    frame-period-summary count;
    symbol-period count;
    }
  protocol-down;
}
```

Syntax: EX Series Switches and NFX Series Devices

```
event {
  link-adjacency-loss;
  link-event-rate {
    frame-error  count;
    frame-period count;
    frame-period-summary  count;
    symbol-period count;
  }
}
```

Hierarchy Level

[edit protocols oam ethernet link-fault-management action-profile]

Release Information

Statement introduced in Junos OS Release 8.5 for MX, M, T, ACX Series routers, SRX Series firewalls and EX Series switches.
Statement introduced in Junos OS Release 9.4 for EX Series switches and NFX devices.

Description

Configure link events in an action profile for Ethernet OAM link fault management (LFM).

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Monitoring Protocol Status*
- *Configuring Ethernet OAM Link Fault Management (CLI Procedure)*

Copyright © 2018, Juniper Networks, Inc.

event-thresholds

Syntax

```
event-thresholds {
   frame-error count;
   frame-period count;
   frame-period-summary count;
   symbol-period count;
}
```

Hierarchy Level [edit protocols oam link-fault-management interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.4.

Description Configure threshold limit values for link events in periodic OAM PDUs.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *Configuring Threshold Values for Local Fault Events on an Interface*

fast-aps-switch

Syntax fast-aps-switch;

Hierarchy Level [edit interfaces *interface-name* sonet-options aps]

Release Information Statement introduced in Junos OS Release 12.1.

Description (M320 routers with Channelized OC3/STM1 Circuit Emulation PIC with SFP only, EX Series switches, and MX series routers with Channelized OC3/STM1 Circuit Emulation PIC with SFP only using container interfaces) Reduce the Automatic Protection Switching (APS) switchover time in Layer 2 circuits.

NOTE:

- The fast APS switching feature is supported only within a single chassis on a MX series router using a container interface.
- Configuring this statement reduces the APS switchover time only when the Layer 2 circuit encapsulation type for the interface receiving traffic from a Layer 2 circuit neighbor is SAToP.
- When the **fast-aps-switch** statement is configured in revertive APS mode, you must configure an appropriate value for revert time to achieve reduction in APS switchover time.
- To prevent the logical interfaces in the data path from being shut down, configure appropriate hold-time values on all the interfaces in the data path that support TDM.
- The **fast-aps-switch** statement cannot be configured when the APS **annex-b** option is configured.
- The interfaces that have the **fast-aps-switch** statement configured cannot be used in virtual private LAN service (VPLS) environments.

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

Related Documentation

- *Reducing APS Switchover Time in Layer 2 Circuits*

 Copyright © 2018, Juniper Networks, Inc.

f-max-period

Syntax f-max-period *number*;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description For all adaptive services interfaces and for ISDN interfaces on J Series Services Routers. Specify the maximum number of compressed packets allowed between the transmission of full headers in a compressed Real-Time Transport Protocol (RTP) traffic stream.

Options *number*—Maximum number of packets. The value can be from 1 through **65535**.

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

Related Documentation
- *Configuring Bandwidth on Demand*
- *Junos OS Services Interfaces Library for Routing Devices*

facility-override

Syntax facility-override *facility-name*;

Hierarchy Level [edit interfaces *interface-name* services-options sysloghost *hostname*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Override default facility for system log reporting.

Options *facility-name*—Name of facility that overrides the default assignment.

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

Related Documentation
- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2018, Juniper Networks, Inc.

failover-delay

Syntax failover-delay *milliseconds*;

Hierarchy Level [edit protocols vrrp]

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure the failover delay for VRRP and VRRP for IPv6 operations.

Options ***milliseconds***—Specify the failover delay time, in milliseconds.

Range: 50 through 2000

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

Related Documentation

- *Configuring VRRP and VRRP for IPv6*

 Copyright © 2018, Juniper Networks, Inc.

family (Dynamic Standard Interface)

Syntax

```
family family {
  access-concentrator name;
  address address;
  direct-connect;
  duplicate-protection;
  dynamic-profile profile-name;
  filter {
    adf {
      counter;
      input-precedence precedence;
      not-mandatory;
      output-precedence precedence;
      rule rule-value;
    }
    input filter-name {
      precedence precedence;
      shared-name filter-shared-name;
    }
    output filter-name {
      precedence precedence;
      shared-name filter-shared-name;
    }
  }
  mac-validate (loose | strict);
  max-sessions number;
  max-sessions-vsa-ignore;
  rpf-check {
    fail-filter filter-name;
    mode loose;
  }
  service {
    input {
      service-set service-set-name {
        service-filter filter-name;
      }
      post-service-filter filter-name;
    }
    output {
      service-set service-set-name {
        service-filter filter-name;
      }
    }
  }
  service-name-table table-name;
  short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
    maximum-seconds> <filter [aci]>;
  unnumbered-address interface-name <preferred-source-address address>;
}
```

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number*]

Copyright © 2018, Juniper Networks, Inc.

Release Information Statement introduced in Junos OS Release 9.2.
pppoe option added in Junos OS Release 11.2.

Description Configure protocol family information for the logical interface.

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

Options *family*—Protocol family:

- **inet**—IP version 4 suite
- **inet6**—IP version 6 suite
- **pppoe**—(MX Series routers with MPCs only) Point-to-Point Protocol over Ethernet
- **vpls**—Virtual private LAN service

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

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

Related Documentation

- *Example: Configuring Static Routing on Logical Systems*
- Configuring the Protocol Family on page 206

 Copyright © 2018, Juniper Networks, Inc.

family

Syntax

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

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


Copyright © 2018, Juniper Networks, Inc.

```
            (preempt | no-preempt) {
              hold-time seconds;
            }
            priority number;
            track {
              interface interface-name {
                bandwidth-threshold bits-per-second priority-cost priority;
                priority-cost priority;
              }
              priority-hold-time seconds;
              route prefix routing-instance instance-name priority-cost priority;
              }
            }
            virtual-address [ addresses ];
            }
          virtual-link-local-address ipv6-address;
        }
      }
```

Hierarchy Level

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

Release Information

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

Description

Configure protocol family information for the logical interface.

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

Options *family*—Protocol family:

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

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

- **bridge**—(M Series and T Series routers only) Configure only when the physical interface is configured with **ethernet-bridge** type encapsulation or when the logical interface is configured with **vlan-bridge** type encapsulation. You can optionally configure this protocol family for the logical interface on which you configure VPLS.
- **ethernet-switching**—(M Series and T Series routers only) Configure only when the physical interface is configured with **ethernet-bridge** type encapsulation or when the logical interface is configured with **vlan-bridge** type encapsulation
- **ccc**—Circuit cross-connect protocol suite. You can configure this protocol family for the logical interface of CCC physical interfaces. When you use this encapsulation type, you can configure the **ccc** family only.
- **inet**—Internet Protocol version 4 suite. You must configure this protocol family for the logical interface to support IP protocol traffic, including Open Shortest Path First (OSPF), Border Gateway Protocol (BGP), Internet Control Message Protocol (ICMP), and Internet Protocol Control Protocol (IPCP).
- **inet6**—Internet Protocol version 6 suite. You must configure this protocol family for the logical interface to support IPv6 protocol traffic, including Routing Information Protocol for IPv6 (RIPng), Intermediate System-to-Intermediate System (IS-IS), BGP, and Virtual Router Redundancy Protocol for IPv6 (VRRP).
- **iso**—International Organization for Standardization Open Systems Interconnection (ISO OSI) protocol suite. You must configure this protocol family for the logical interface to support IS-IS traffic.
- **mlfr-end-to-end**—Multilink Frame Relay FRF.15. You must configure this protocol or multilink Point-to-Point Protocol (MLPPP) for the logical interface to support multilink bundling.
- **mlfr-uni-nni**—Multilink Frame Relay FRF.16. You must configure this protocol or **mlfr-end-to-end** for the logical interface to support link services and voice services bundling.
- **multilink-ppp**—Multilink Point-to-Point Protocol. You must configure this protocol (or **mlfr-end-to-end**) for the logical interface to support multilink bundling.
- **mpls**—Multiprotocol Label Switching (MPLS). You must configure this protocol family for the logical interface to participate in an MPLS path.
- **pppoe**—Point-to-Point Protocol over Ethernet
- **tcc**—Translational cross-connect protocol suite. You can configure this protocol family for the logical interface of TCC physical interfaces.

 Copyright © 2018, Juniper Networks, Inc.

- **tnp**—Trivial Network Protocol. This protocol is 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, as discussed in "Understanding Internal Ethernet Interfaces" on page 15.
- **vpls**—(M Series and T Series routers only) Virtual private LAN service. You can optionally configure this protocol family for the logical interface on which you configure VPLS. VPLS provides an Ethernet-based point-to-multipoint Layer 2 VPN to connect customer edge (CE) routers across an MPLS backbone. When you configure a VPLS encapsulation type, the **family vpls** statement is assumed by default.

  MX Series routers support dynamic profiles for VPLS pseudowires, VLAN identifier translation, and automatic bridge domain configuration.

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

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

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

Related Documentation

- Configuring the Protocol Family on page 206

fastether-options

Syntax

```
fastether-options {
  802.3ad {
    aex (primary | backup);
    lacp {
      port-priority;
    }
  }
  (flow-control | no-flow-control);
  ignore-l3-incompletes;
  ingress-rate-limit rate;
  (loopback | no-loopback);
  mpls {
    pop-all-labels {
      required-depth number;
    }
  }
  source-address-filter {
    mac-address;
  }
  (source-filtering | no-source-filtering);
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure Fast Ethernet-specific interface properties.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *Ethernet Interfaces Overview*

 Copyright © 2018, Juniper Networks, Inc.

fcs

Syntax fcs (16 | 32);

Hierarchy Level
[edit interfaces e1-*fpc/pic/port*],
[edit interfaces t1-*fpc/pic/port*],
[edit interfaces *interface-name* ds0-options],
[edit interfaces *interface-name* e1-options],
[edit interfaces *interface-name* e3-options],
[edit interfaces *interface-name* sonet-options],
[edit interfaces *interface-name* t1-options],
[edit interfaces *interface-name* t3-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers.

Description For E1/E3, SONET/SDH, and T1/T3 interfaces, configure the frame checksum (FCS) on the interface. The checksum must be the same on both ends of the interface.

On a channelized OC12 interface, the SONET/SDH **fcs** statement is not supported. To configure FCS on each DS3 channel, you must include the **t3-options fcs** statement in the configuration for each channel. For SONET/SDH, the channelized OC12 interface supports DS3 to STS-1 to OC12. For SDH, the channelized OC12 interface supports *N*xDS3 to *N*xVC3 to AU3 to STM.

NOTE: **When configuring E1 or T1 interfaces on 10-port Channelized E1/T1 IQE PICs, the fcs statement must be included at the [edit interfaces e1-*fpc/pic/port*] or [edit interfaces t1-*fpc/pic/port*] hierarchy level as appropriate.**

Options **16**—Use a 16-bit frame checksum on the interface.

32—Use a 32-bit frame checksum on the interface. Using a 32-bit checksum provides more reliable packet verification, but some older equipment might not support 32-bit checksums.

Default: 16

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

Related Documentation
- *Configuring the E1 Frame Checksum*
- *Configuring the E3 Frame Checksum*
- *Configuring the SONET/SDH Frame Checksum*
- *Configuring the T1 Frame Checksum*

Copyright © 2018, Juniper Networks, Inc.

- *Configuring the T3 Frame Checksum*

feac-loop-respond

Syntax (feac-loop-respond | no-feac-loop-respond);

Hierarchy Level [edit interfaces *interface-name* t3-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For T3 interfaces only, configure the router so a remote CSU can place the local router into loopback.

If you configure remote or local loopback with the T3 **loopback** statement, the router does not respond to FEAC requests from the CSU even if you include the **feac-loop-respond** statement in the configuration. For the router to respond, you must delete the **loopback** statement from the configuration.

You must rollback the setting done on the remote CSU prior to deactivating the **feac-loop-respond** statement. If the remote CSU cannot comply, clear the remote loop through local configuration to achieve the cleanup. For example, configure remote loopback on the interface and then delete the remote loopback.

Default The router does not respond to FEAC requests.

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

Related Documentation

- *Configuring the T3 FEAC Response*
- loopback (ADSL, DS0, E1/E3, SONET/SDH, SHDSL, and T1/T3) on page 751
- remote-loopback-respond on page 928

 Copyright © 2018, Juniper Networks, Inc.

fec (gigether)

Syntax fec (fec91 | none)

Hierarchy Level [edit interfaces *interface-name* gigether-options]

Release Information Statement introduced in Junos OS Release 16.1R1
Statement introduced in Junos OS Release 16.1X65D30 for PTX1000 routers
Statement introduced in Junos OS Release 17.1R1 for PTX5000 routers

Description (MX Series Routers with MPC7E, MPC8E, and MPC9E, PTX1000, PTX5000) Enable or disable RS-FEC (Reed-Solomon Forward Error Correction) for a 100-Gigabit Ethernet interface. By default, the Junos OS software enables or disables forward error correction based on the plugged-in optics. For instance, Junos OS software enables RS-FEC for 100G SR4 optics and disables RS-FEC for 100G LR4 optics.

This statement allows you to override the default behavior and explicitly enable or disable RS-FEC. For instance, you can extend the reach of 100G LR4 optics when you explicitly enable RS-FEC for the optics. RS-FEC is compliant with IEEE 802.3-2015 Clause 91.

Once you enable or disable RS-FEC using this statement, this behavior applies to any 100-Gigabit Ethernet optical transceiver installed in the port associated with the interface. Delete the statement and commit the configuration to return to the default behavior.

Default Junos OS software automatically enables or disables RS-FEC based on the type of pluggable optics used.

Options **fec91**—Enables RS-FEC. RS-FEC is compliant with IEEE 802.3-2015 Clause 91.

none—Disables RS-FEC.

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

Related Documentation

- *MPC7E (Multi-Rate) on MX Series Routers Overview*
- *MPC8E on MX Series Routers Overview*
- *MPC9E on MX Series Routers Overview*
- *Determining Transceiver Support for the PTX1000*

Copyright © 2018, Juniper Networks, Inc.

filter

Syntax

```
filter {
  group filter-group-number;
  input filter-name;
  input-list [ filter-names ];
  output filter-name;
  output-list [ filter-names ];
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description

NOTE: On EX Series switches, the **group**, **input-list**, **output-filter** statements are not supported under the **[edit interfaces *interface-name* unit *logical-unit-number* family inet]**, **[edit interfaces *interface-name* unit *logical-unit-number* family inet6]**, and **[edit interfaces *interface-name* unit *logical-unit-number* family mpls]** hierarchies.

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.

Options

group *filter-group-number*—Define an interface to be part of a filter group. The default filter group number is 0.

Range: 0 through 255

input *filter-name*—Name of one filter to evaluate when packets are received on the interface.

output *filter-name*—Name of one filter to evaluate when packets are transmitted on the interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- Applying a Filter to an Interface on page 235
- *Junos OS Administration Library*

Copyright © 2018, Juniper Networks, Inc.

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

filter (Applying to an Interface)

Syntax

```
filter {
  input filter-name;
  output filter-name;
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Apply a filter to an interface. You can also use filters for encrypted traffic. When you configure filters, you can configure the family **inet**, **inet6**, **mpls**, or **vpls** only.

Options

input *filter-name*—Name of one filter to evaluate when packets are received on the interface.

output *filter-name*—Name of one filter to evaluate when packets are transmitted 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

- *simple-filter*
- *Configuring and Applying Tricolor Marking Policers*
- *Example: Classifying Packets Based on Their Destination Address*
- *Example: Configuring and Verifying a Complex Multifield Filter*
- *Example: Writing Different DSCP and EXP Values in MPLS-Tagged IP Packets*
- *Configuring a Simple Filter*
- *Configuring Policers Based on Logical Interface Bandwidth*
- *Effect of Two-Color Policers on Shaping Rate Changes*

flexible-vlan-tagging

Syntax

flexible-vlan-tagging;

Hierarchy Level

[edit interfaces aex],
[edit interfaces ge-*fpc/pic/port*],
[edit interfaces et-*fpc/pic/port*],
[edit interfaces ps0],
[edit interfaces xe-*fpc/pic/port*]

Release Information

Statement introduced in Junos OS Release 8.1.
Support for aggregated Ethernet added in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers.
Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.

Description

Support simultaneous transmission of 802.1Q VLAN single-tag and dual-tag frames on logical interfaces on the same Ethernet port, and on pseudowire logical interfaces.

This statement is supported on M Series and T Series routers, for Fast Ethernet and Gigabit Ethernet interfaces only on Gigabit Ethernet IQ2 and IQ2-E, IQ, and IQE PICs, and for aggregated Ethernet interfaces with member links in IQ2, IQ2-E, and IQ PICs or in MX Series DPCs, or on Ethernet interfaces for PTX Series Packet Transport Routers or 100-Gigabit Ethernet Type 5 PIC with CFP.

This statement is supported on Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces on EX Series and QFX Series switches.

Required Privilege Level

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

Related Documentation

- *Enabling VLAN Tagging*
- *Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers*
- *Configuring Double-Tagged VLANs on Layer 3 Logical Interfaces*

Copyright © 2018, Juniper Networks, Inc.

flow-control

Syntax (flow-control | no-flow-control);

Hierarchy Level [edit interfaces *interface-name* aggregated-ether-options],
[edit interfaces *interface-name* ether-options],
[edit interfaces *interface-name* fastether-options],
[edit interfaces *interface-name* gigether-options],
[edit interfaces *interface-name* multiservice-options],
[edit interfaces interface-range *name* aggregated-ether-options],
[edit interfaces interface-range *name* ether-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 in EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description For aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only, explicitly enable flow control, which regulates the flow of packets from the router or switch to the remote side of the connection. Enabling flow control is useful when the remote device is a Gigabit Ethernet switch. Flow control is not supported on the 4-port Fast Ethernet PIC.

NOTE: On the Type 5 FPC, to prioritize control packets in case of ingress oversubscription, you must ensure that the neighboring peers support MAC flow control. If the peers do not support MAC flow control, then you must disable flow control.

Default Flow control is enabled.

NOTE: Flow control is enabled by default only on physical interfaces and it is disabled by default on aggregated Ethernet interfaces.

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

Related Documentation

- *Configuring Flow Control*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure) for EX Series Switches with ELS support*

Copyright © 2018, Juniper Networks, Inc.

flow-control-options

Syntax

```
flow-control-options {
  down-on-flow-control;
  dump-on-flow-control;
  reset-on-flow-control;
  up-on-flow-control;
}
```

Hierarchy Level [edit interfaces mo-*fpc/pic/port* multiservice-options]

Release Information Statement introduced before Junos OS Release 8.4.

Description Configure the flow control options for application recovery in case of a prolonged flow control failure.

- **down-on-flow-control**—Bring interface down during prolonged flow control.
- **dump-on-flow-control**—Cause core dump during prolonged flow control.

NOTE: Starting with Junos OS Release 15.1, on MX Series routers with MS-MICs and MS-MPCs, instead of an eJunos kernel core file, the multiservices PIC management daemon (mspmand) core file is generated when a prolonged flow control failure occurs and when you configure the setting to generate a core dump during prolonged flow control (by using the **dump-on-flow-control** option with the **flow-control-options** statement). The watchdog functionality continues to generate a kernel core file in such scenarios.

- **reset-on-flow-control**—Reset interface during prolonged flow control.

NOTE: Starting in Junos OS Release 16.1R7, the **reset-on-flow-control** option has no effect on the MS-MIC, MS-MPC, MS-DPC, MS-PIC 100, MS-PIC 400, and MS-PIC 500 line cards. This is because starting in Release 16.1R7, Junos OS restarts these line cards to recover them from stuck state due to prolonged flow control.

- **up-on-flow-control**—Cause interface to remain in stuck state until you manually restart the PICs.

NOTE: Starting in Junos OS Release 16.1R7, if interfaces on an MS-PIC or MS-DPC are in stuck state because of prolonged flow control, Junos OS restarts the service PICs to recover them from this state. However, if you want the PICs to remain in stuck state until you manually restart the PICs,

Copyright © 2018, Juniper Networks, Inc.

configure the **up-on-flow-control** option. In releases before Release 16.1R7, there is no action taken to recover service PICs from this state unless one of the options for the **flow-control-options** statement is configured, or service PIC is manually restarted.

Usage Guidelines See *Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches*.

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

force

Syntax force (protect | working);

Hierarchy Level [edit interfaces *interface-name* sonet-options aps]

Release Information Statement introduced before Junos OS Release 7.4.

Description Perform a forced switch between the protect and working circuits. This statement is honored only if there are no higher-priority reasons to switch. It can be overridden by a signal failure on the protect circuit, thus causing a switch to the working circuit.

Options **protect**—Request the circuit to become the protect circuit.

working—Request the circuit to become the working circuit.

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

Related Documentation

- *Configuring Switching Between the Working and Protect Circuits*
- request on page 930

forward-and-send-to-re

Syntax forward-and-send-to-re;

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

- Configuring Targeted Broadcast on page 272
- **targeted-broadcast on page 1010**
- Understanding Targeted Broadcast on page 271

 Copyright © 2018, Juniper Networks, Inc.

forwarding-class (ATM2 IQ Scheduler Maps)

Syntax

```
forwarding-class class-name {
  epd-threshold cells plp1 cells;
  linear-red-profile profile-name;
  priority (high | low);
  transmit-weight (cells number | percent number);
}
```

Hierarchy Level [edit interfaces at-*fpc/pic/port* atm-options scheduler-maps *map-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For ATM2 IQ interfaces only, define forwarding class name and option values.

Options ***class-name***—Name of forwarding class.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *ATM2 IQ VC Tunnel CoS Components Overview*
- *Applying Scheduler Maps to ATM Interfaces*

Copyright © 2018, Juniper Networks, Inc.

forwarding-class (Gigabit Ethernet IQ Classifier)

Syntax

forwarding-class *class-name* {
  loss-priority (high | low);
}

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile
  ethernet-policer-profile output-priority-map classifier premium]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

For Gigabit Ethernet IQ interfaces only, define forwarding class name and option values.

Options

class-name—Name of forwarding class.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Specifying an Output Priority Map*
- input-priority-map on page 677
- **forwarding-class** statement in the *Class of Service Feature Guide for Routing Devices and EX9200 Switches*

 Copyright © 2018, Juniper Networks, Inc.

forward-only

Syntax forward-only;

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

- Configuring Targeted Broadcast on page 272
- **targeted-broadcast on page 1010**
- Understanding Targeted Broadcast on page 271

fragment-threshold

Syntax fragment-threshold *bytes*;

Hierarchy Level [edit interfaces *interface-name* mlfr-uni-nni-bundle-options],
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For multilink, link services, and voice services interfaces, set the fragmentation threshold.

Options *bytes*—Maximum size, in bytes, for multilink packet fragments. Any nonzero value must be a multiple of 64 bytes.

Range: 128 through 16,320 bytes

Default: 0 bytes (no fragmentation)

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

frame-error

Syntax frame-error *count*;

Hierarchy Level [edit protocols oam ethernet link-fault-management action-profile event link-event-rate], [edit protocols oam link-fault-management interface *interface-name* event-thresholds]

Release Information Statement introduced in Junos OS Release 8.4.

Description Threshold for sending frame error events or taking the action specified in the action profile.

A frame error is any frame error on the underlying physical layer. The threshold is reached when the number of frame errors reaches the configured value within the window.

The window or period during which frame errors are counted is 5 seconds or multiples of it (with a maximum value of 1 minute). This window denotes the duration as intervals of 100 milliseconds, encoded as a 16-bit unsigned integer. This window is not configurable in Junos OS. According to the IEEE 802.3ah standard, the default value of the frame-errors window is 1 second. This window has a lower bound of 1 second and an upper bound of 1 minute.

Options ***count***—Threshold count for frame error events.

Range: 0 through 100

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

Related Documentation

- *Configuring Threshold Values for Local Fault Events on an Interface*
- *Configuring Threshold Values for Fault Events in an Action Profile*

 Copyright © 2018, Juniper Networks, Inc.

frame-period

Syntax frame-period *count*;

Hierarchy Level [edit protocols oam ethernet link-fault-management action-profile event link-event-rate], [edit protocols oam link-fault-management interface *interface-name* event-thresholds]

Release Information Statement introduced in Junos OS Release 8.4.

Description Threshold for sending frame period error events or taking the action specified in the action profile.

A frame error is any frame error on the underlying physical layer. The frame period threshold is reached when the number of frame errors reaches the configured value within the period window. The default period window is the number of minimum-size frames that can be transmitted on the underlying physical layer in 1 second. The window is not configurable.

Options ***count***—Threshold count for frame period error events.

Range: 0 through 100

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

Related Documentation

- *Configuring Threshold Values for Local Fault Events on an Interface*
- *Configuring Threshold Values for Fault Events in an Action Profile*

Copyright © 2018, Juniper Networks, Inc.

frame-period-summary

Syntax

frame-period-summary *count*;

Hierarchy Level

[edit protocols oam ethernet link-fault-management action-profile event link-event-rate],
[edit protocols oam link-fault-management interface *interface-name* event-thresholds]

Release Information

Statement introduced in Junos OS Release 8.4.

Description

Threshold for sending frame period summary error events or taking the action specified in the action profile.

An errored frame second is any 1-second period that has at least one errored frame. This event is generated if the number of errored frame seconds is equal to or greater than the specified threshold for that period window. The default window is 60 seconds. The window is not configurable.

Options

count—Threshold count for frame period summary error events.

Range: 0 through 100

Required Privilege Level

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

Related Documentation

- *Configuring Threshold Values for Local Fault Events on an Interface*
- *Configuring Threshold Values for Fault Events in an Action Profile*

 Copyright © 2018, Juniper Networks, Inc.

framing (E1, E3, and T1 Interfaces)

Syntax framing (g704 | g704-no-crc4 | g.751 | g.832 | unframed | sf | esf);

Hierarchy Level [edit interfaces ce1-*fpc/pic/port*],
[edit interfaces ct1-*fpc/pic/port*],
[edit interfaces at-*fpc/pic/port* e3-options],
[edit interfaces e1-*fpc/pic/port* e1-options],
[edit interfaces t1-*fpc/pic/port* t1-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers.

Description Configure the framing format.

NOTE: **When configuring CE1 or CT1 interfaces on 10-port Channelized E1/T1 IQE PICs, the framing statement must be included at the [edit interfaces ce1-*fpc/pic/port*] or [edit interfaces ct1-*fpc/pic/port*] hierarchy level as appropriate.**

Default **esf** for T1 interfaces; **g704** for E1 interfaces. There is no default value for E3 over ATM interfaces.

Options **esf**—Extended superframe (ESF) mode for T1 interfaces.

g704—G.704 framing format for E1 interfaces.

g704-no-crc4—G.704 framing with no cyclic redundancy check 4 (CRC4) for E1 interfaces.

g.751—G.751 framing format for E3 over ATM interfaces.

g.832—G.832 framing format for E3 over ATM interfaces.

sf—Superframe (SF) mode for T1 interfaces.

unframed—Unframed mode for E1 interfaces.

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

Related Documentation

- *Configuring E1 Framing*
- *Configuring E3 and T3 Parameters on ATM Interfaces*
- *Configuring T1 Framing*

Copyright © 2018, Juniper Networks, Inc.

framing (10-Gigabit Ethernet Interfaces)

Syntax framing (lan-phy | wan-phy);

Hierarchy Level [edit interfaces xe-*fpc/pic/port*]

[edit interfaces et-*fpc/pic/port*] (PTX Series Packet Transport Routers and MX Series Routers)

Release Information Statement introduced in Junos OS Release 8.0.
Statement introduced in Junos OS Release 12.3R2 for PTX Series Packet Transport Routers.

Description For routers supporting the 10-Gigabit Ethernet interface, configure the framing format. WAN PHY mode is supported on MX240, MX480, MX960, T640, T1600,T4000, and PTX Series Packet Transport Routers routers only.

NOTE:

- The T4000 Core Router supports only LAN PHY mode in Junos OS Release 12.1R1. Starting with Junos OS Release 12.1R2, WAN PHY mode is supported on the T4000 routers with the 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP). Starting with Junos OS Release 12.2, WAN PHY mode is supported on the T4000 routers with the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-24XGE-SFPP).
- On PTX Series routers, WAN PHY mode is supported only on the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ .
- When the PHY mode changes, interface traffic is disrupted because of port reinitialization.

Default Operates in LAN PHY mode.

Options **lan-phy**—10GBASE-R interface framing format that bypasses the WIS sublayer to directly stream block-encoded Ethernet frames on a 10-Gigabit Ethernet serial interface.

wan-phy—10GBASE-W interface framing format that allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and SONET devices.

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

Related Documentation

- *10-Gigabit Ethernet Framing Overview*
- *Configuring SONET Options for 10-Gigabit Ethernet Interfaces*

 Copyright © 2018, Juniper Networks, Inc.

framing (SONET and SDH Interfaces)

Syntax framing (sdh | sonet);

Hierarchy Level [edit interfaces so-*fpc/pic/port*]

Release Information Statement introduced in Junos OS Release 8.1.

Description This functionality allows you to mix SONET and SDH modes on interfaces on the same PIC.

- For the 4-port OC48 PIC with SFP installed and the 4-port OC192 PIC in T Series and M Series routers, configure SONET or SDH framing on a per-port basis.
- For 1-port OC192/STM64 MICs with XFP on MX Series routers, configure the SONET or SDH framing on the single port.

Default Default framing mode is **SONET**.

Options **sdh**—SDH framing.

sonet—SONET framing.

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

Related Documentation

- *Configuring SONET/SDH Framing Mode for Ports*

gigether-options

Syntax

```
gigether-options {
  802.3ad {
    aex (primary | backup);
    lacp {
      port-priority;
    }
  }
  (asynchronous-notification | no-asynchronous-notification);
  (auto-negotiation | no-auto-negotiation) remote-fault <local-interface-online |
    local-interface-offline>;
  fec
  (flow-control | no-flow-control);
  ignore-l3-incompletes;
  (loopback | no-loopback);
  mpls {
    pop-all-labels {
      required-depth number;
    }
  }
  no-auto-mdix
  source-address-filter {
    mac-address;
  }
  (source-filtering | no-source-filtering);
  speed
  ethernet-switch-profile {
    (mac-learn-enable | no-mac-learn-enable);
    tag-protocol-id [ tpids ];
    ethernet-policer-profile {
      input-priority-map {
        ieee802.1p premium [ values ];
      }
      output-priority-map {
        classifier {
          premium {
            forwarding-class class-name {
              loss-priority (high | low);
            }
          }
        }
      }
      policer cos-policer-name {
        aggregate {
          bandwidth-limit bps;
          burst-size-limit bytes;
        }
        premium {
          bandwidth-limit bps;
          burst-size-limit bytes;
        }
      }
    }
```

 Copyright © 2018, Juniper Networks, Inc.

```
    }
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure Gigabit Ethernet specific interface properties.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *Ethernet Interfaces Overview*
- *gigether-options (ACX Series)*

gratuitous-arp-reply

Syntax (gratuitous-arp-reply | no-gratuitous-arp-reply);

Hierarchy Level [edit interfaces *interface-name*]
[edit interfaces **interface-range** *interface-range-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 in EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description For Ethernet interfaces, enable updating of the Address Resolution Protocol (ARP) cache for gratuitous ARPs.

Default Updating of the ARP cache is disabled on all Ethernet interfaces.

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

Related Documentation

- *Configuring Gratuitous ARP*
- no-gratuitous-arp-request on page 816

guard-interval

Syntax

guard-interval *number*;

Hierarchy Level

[edit protocols protection-group ethernet-ring *ring-name*]

Release Information

Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

Description

When a link goes down, the ring protection link (RPL) activates. When the downed link comes back up, the RPL link receives notification, restores the link, and waits for the restore interval before issuing another block on the same link. This configuration is a global configuration and applies to all Ethernet rings if the Ethernet ring does not have a more specific configuration for this value. If no parameter is configured at the protection group level, the global configuration of this parameter uses the default value.

Options

number—Guard timer interval, in milliseconds.

Range: 10 through 2000 ms

Default: 500 ms

Required Privilege Level

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

Related Documentation

- *Ethernet Ring Protection Switching Overview*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

Copyright © 2018, Juniper Networks, Inc.

hardware-assisted-timestamping

Syntax hardware-assisted-timestamping;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management performance-monitoring]

Release Information Statement introduced in Junos OS Release 9.5.

Description For Ethernet interfaces on Enhanced and Enhanced Queuing Dense Port Concentrators (DPCs) in MX Series routers only, enable hardware-assisted timestamping support for Ethernet frame delay measurement.

By default, the ETH-DM feature calculates frame delays using software-based timestamping of the ETH-DM PDU frames sent and received by the MEPs in the session. As an option that can increase the accuracy of ETH-DM calculations when the DPC is loaded with heavy traffic in the receive direction, you can enable hardware-assisted timestamping of session frames in the receive direction.

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

Related Documentation

- *Ethernet Frame Delay Measurements Overview*
- *Guidelines for Configuring Routers to Support an ETH-DM Session*
- *Enabling the Hardware-Assisted Timestamping Option*

Copyright © 2018, Juniper Networks, Inc.

high-plp-threshold

Syntax high-plp-threshold *percent*;

Hierarchy Level [edit interfaces at-*fpc/pic/port* atm-options linear-red-profiles *profile-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For ATM2 IQ interfaces only, define CoS VC drop profile fill-level percentage when linear RED is applied to cells with high PLP. When the fill level exceeds the defined percentage, packets with high PLP are randomly dropped by RED. This statement is mandatory.

Options *percent*—Fill-level percentage when linear RED is applied to cells with PLP.

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

Related Documentation

- *ATM2 IQ VC Tunnel CoS Components Overview*
- high-plp-max-threshold on page 640
- low-plp-max-threshold on page 758
- low-plp-threshold on page 759
- queue-depth on page 914

 Copyright © 2018, Juniper Networks, Inc.

hello-timer

Syntax hello-timer *milliseconds*;

Hierarchy Level [edit interfaces *interface-name* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services and voice services interfaces only, configure the rate at which hello messages are sent. A hello message is transmitted after a period defined in milliseconds has elapsed.

Options **milliseconds**—The rate at which hello messages are sent.

Range: 1 through 180 milliseconds

Default: 10 milliseconds

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*
- acknowledge-timer on page 417
- address on page 424

hierarchical-policer

Syntax

```
hierarchical-policer name {
aggregate {
  if-exceeding {
    bandwidth-limit bandwidth;
    burst-size-limit burst;
  }
  then {
    discard;
  }
}
premium {
  if-exceeding {
    bandwidth-limit bandwidth;
    burst-size-limit burst;
  }
  then {
    discard;
  }
}
```

Hierarchy Level [edit firewall]

Release Information Statement introduced in Junos OS Release 9.5.

Description For M40e, M120, and M320 (with FFPC and SFPC) edge routers and T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs, specify a hierarchical policer.

Options Options are described separately.

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

Related Documentation

- Applying Policers on page 226
- *Class of Service Feature Guide for Routing Devices and EX9200 Switches*

 Copyright © 2018, Juniper Networks, Inc.

hierarchical-scheduler (Subscriber Interfaces on MX Series Routers)

Syntax

```
hierarchical-scheduler {
  implicit-hierarchy;
  maximum-hierarchy-levels number;
}
```

Hierarchy Level

[edit interfaces *interface-name*]

Release Information

Statement introduced in Junos OS Release 10.1.
implicit-hierarchy option added in Junos OS Release 13.1.
Support on GRE tunnel interfaces configured on physical interfaces on MICs or MPCs in MX Series routers added in Junos OS Release 13.3.
Support for up to four hierarchy levels added in Junos OS Release 16.1.

Description

Configure hierarchical scheduling options on the interface.

The statement is supported on the following interfaces:

- MIC and MPC interfaces in MX Series routers
- GRE tunnel interfaces configured on physical interfaces hosted on MIC or MPC line cards in MX Series routers

To enable hierarchical scheduling on MX Series routers, configure the **hierarchical-scheduler** statement at each member physical interface level of a particular aggregated Ethernet interface as well as at that aggregated Ethernet interface level. On other routing platforms, it is enough if you include this statement at the aggregated Ethernet interface level.

Options

implicit-hierarchy—Configure four-level hierarchical scheduling. When you include the **implicit-hierarchy** option, a hierarchical relationship is formed between the CoS scheduler nodes at level 1, level 2, level 3, and level 4. The **implicit-hierarchy** option is supported only on MPC/MIC subscriber interfaces and interface sets on MX Series routers.

maximum-hierarchy-levels ***number***—Specify the maximum number of hierarchical scheduling levels allowed for node scaling, from 2 through 4 levels. The default number of levels is 3. The **maximum-hierarchy-levels** option is supported on MPC/MIC or EQ DPC subscriber interfaces and interface sets on MX Series routers.

- If you set **maximum-hierarchy-levels** to **2**, interface sets are not allowed. In this case, if you configure a level 2 interface set, you generate Packet Forwarding Engine errors.
- If you do not include the **maximum-hierarchy-levels** option, keeping the default number of hierarchy levels at 3, interface sets can be at either level 2 or level 3, depending on whether the member logical interfaces within the interface set have a traffic control profile. If any member logical interface has a traffic control profile, then the interface set is a level 2 CoS scheduler node. If no member logical interface has a traffic control profile, the interface set is at level 3.

Copyright © 2018, Juniper Networks, Inc.

CAUTION: *MPC3E*, *32x10GE MPC4E*, and *2x100GE + 8x10GE MPC4E* MPCs support only two levels of scheduling hierarchy. When enabling hierarchical scheduling on these cards, you must explicitly set **maximum-hierarchy-levels** to 2.

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

Related Documentation

- *Understanding Hierarchical CoS for Subscriber Interfaces*
- *Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links*
- *Configuring Hierarchical Schedulers for CoS*
- *Configuring Hierarchical CoS on a Static PPPoE Subscriber Interface*
- *Hierarchical CoS on MPLS Pseudowire Subscriber Interfaces Overview*

high-plp-max-threshold

Syntax high-plp-max-threshold *percent*;

Hierarchy Level [edit interfaces at-*fpc/pic/port* atm-options linear-red-profiles *profile-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For ATM2 IQ interfaces only, define the drop profile fill-level for the high PLP CoS VC. When the fill level exceeds the defined percentage, all packets are dropped.

Options *percent*—Fill-level percentage when linear random early discard (RED) is applied to cells with PLP.

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

Related Documentation

- *ATM2 IQ VC Tunnel CoS Components Overview*
- low-plp-max-threshold on page 758
- low-plp-threshold on page 759
- queue-depth on page 914

Copyright © 2018, Juniper Networks, Inc.

high-plp-threshold

Syntax high-plp-threshold *percent*;

Hierarchy Level [edit interfaces at-*fpc/pic/port* atm-options linear-red-profiles *profile-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For ATM2 IQ interfaces only, define CoS VC drop profile fill-level percentage when linear RED is applied to cells with high PLP. When the fill level exceeds the defined percentage, packets with high PLP are randomly dropped by RED. This statement is mandatory.

Options ***percent***—Fill-level percentage when linear RED is applied to cells with PLP.

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

Related Documentation

- *ATM2 IQ VC Tunnel CoS Components Overview*
- high-plp-max-threshold on page 640
- low-plp-max-threshold on page 758
- low-plp-threshold on page 759
- queue-depth on page 914

hold-interval (OAM)

Syntax hold-interval *minutes*;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* continuity-check]

Release Information Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

Description The time to wait in minutes before flushing the maintenance association end point (MEP) database, if no updates occur. The configurable range is 1 minute through 30240 minutes.The default value is 10 minutes.

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

Options *minutes*—Time to wait, in minutes.

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

Related Documentation

- *Continuity Check Protocol Parameters Overview*
- *Configuring Continuity Check Protocol Parameters for Fault Detection*

 Copyright © 2018, Juniper Networks, Inc.

hold-interval (Protection Group)

Syntax hold-interval *number*;

Hierarchy Level [edit protocols protection-group ethernet-ring *name*]

Release Information Statement introduced in Junos OS Release 9.4.

Description Specify the hold-off timer interval *for all rings* in 100 millisecond (ms) increments.

Options *number*—Hold-timer interval, in milliseconds.

Range: 0 through 10,000 ms

Default: 100 ms

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

Related Documentation

- *Ethernet Ring Protection Switching Overview*
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*

hold-time (APS)

Syntax hold-time *milliseconds*;

Hierarchy Level [edit interfaces *interface-name* sonet-options aps]

Release Information Statement introduced before Junos OS Release 7.4.

Description Hold-time value to use to determine whether a neighbor APS router is operational.

Options *milliseconds*—Hold-time value.

Range: 1 through 65,534 milliseconds

Default: 3000 milliseconds (3 times the advertisement interval)

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

Related Documentation

- *Configuring APS Timers*
- advertise-interval on page 426

Copyright © 2018, Juniper Networks, Inc.

hold-time (Physical Interface)

Syntax hold-time up *milliseconds* down *milliseconds*;

Hierarchy Level [edit interfaces *interface-name*],
[edit interfaces interface-range *interface-range-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 10.4R5 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Statement introduced in Junos OS Release 12.1 for the SRX Series.

Description Specify the **hold-time** value to use to damp shorter interface transitions milliseconds. The hold timer enables interface damping by not advertising interface transitions until the hold timer duration has passed. When a hold-down timer is configured and the interface goes from up to down, the down hold-time timer is triggered. Every interface transition that occurs during the hold-time is ignored. When the timer expires and the interface state is still down, then the router begins to advertise the interface as being down. Similarly, when a hold-up timer is configured and an interface goes from down to up, the up hold-time timer is triggered. Every interface transition that occurs during the hold-time is ignored. When the timer expires and the interface state is still up, then the router begins to advertise the interface as being up.

NOTE:

- We recommend that you configure the **hold-time** value after determining an appropriate value by performing repeated tests in the actual hardware environment. This is because the appropriate value for **hold-time** depends on the hardware (XFP, SFP, SR, ER, or LR) used in the networking environment.
- The **hold-time** option is not available for controller interfaces.

NOTE: On MX Series routers with MPC3E and MPC4E, we recommend that you do not configure the hold-down timer to be less than 1 second. On MX Series routers with MPC5EQ-100G10G (MPC5EQ) or MPC6E (MX2K-MPC6E) with 100-Gigabit Ethernet MIC with CFP2 OTN interfaces, we recommend that you do not configure the hold-down timer to be less than 3 seconds.

Default Interface transitions are not damped.

 Copyright © 2018, Juniper Networks, Inc.

Options

down ***milliseconds***—Hold time to use when an interface transitions from up to down. Junos OS advertises the transition within 100 milliseconds of the time value you specify.

Range: 0 through 4,294,967,295

Default: 0 (interface transitions are not damped)

up ***milliseconds***—Hold time to use when an interface transitions from down to up. Junos OS advertises the transition within 100 milliseconds of the time value you specify.

Range: 0 through 4,294,967,295

Default: 0 (interface transitions are not damped)

Required Privilege Level

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

Related Documentation

- **advertise-interval on page 426**
- *interfaces (EX Series switches)*
- Physical Interface Damping Overview on page 159
- Damping Shorter Physical Interface Transitions on page 165
- Damping Longer Physical Interface Transitions on page 166

hold-time (SONET/SDH Defect Triggers)

Syntax hold-time up *milliseconds* down *milliseconds*;

Hierarchy Level [edit interfaces *interface-name* sonet-options trigger *defect*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For ATM over SONET/SDH and SONET/SDH interfaces only, apply up and down hold times to SONET/SDH defect triggers. When you apply a down hold time to a defect, the defect must remain present for at least the hold-time period before the interface is marked down. When you apply an up hold time to a defect, the defect must remain absent for at least the hold-time period before the interface is marked up, assuming no other defect is outstanding.

NOTE:

- When up or down hold times are applied to SONET defect triggers of a 10-Gigabit Ethernet WAN-PHY interface, only the defects generated in the WAN Interface Sublayer (WIS) are damped. Therefore, if the hold times are applied to SONET defect triggers only, a 10-Gigabit Ethernet WAN-PHY interface might be marked up or down because of the faults that are generated in other layers, such as the Physical Coding Sublayer (PCS) or Physical Medium Attachment Sublayer (PMA), 10 Gigabit Media Independent Interface (XGMII) Extender Sublayer (XGXS), and Media Access Control (MAC). To damp the interface up or down events of a 10-Gigabit Ethernet WAN-PHY interface, you need to apply up or down hold-times for the interface at the **[edit interfaces *interface-name*]** hierarchy level.
- On M Series and T Series platforms with Channelized SONET IQ PICs and Channelized SONET IQE PICs, the SONET defect alarm trigger **hold-time** statement is not supported.

Default If you do not include this statement, when a defect is detected the interface is marked down immediately, and when the defect becomes absent the interface is marked up immediately.

Options **down *milliseconds***—Hold time to wait before the interface is marked down.

Range: 1 through 65,534 milliseconds

Default: No hold time

up *milliseconds*—Hold time to wait before the interface is marked up.

Range: 1 through 65,534 milliseconds

Default: No hold time

 Copyright © 2018, Juniper Networks, Inc.

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

Related Documentation

- *Configuring SONET/SDH Defect Triggers*
- hold-time (Physical Interface) on page 644

host (Interfaces)

Syntax

```
host hostname {
    services severity-level;
    facility-override facility-name;
    log-prefix prefix-value;
    port port-number;
}
```

Hierarchy Level [edit interfaces *interface-name* services-options syslog]

Release Information Statement introduced before Junos OS Release 7.4.
You can configure multiple system log hosts from Junos OS Release 17.4R1 onwards.

Description Specify the hostname for the system logging utility.

Starting with Junos OS release 17.4R1, you can configure up to a maximum of four system log servers (combination of local system log hosts and remote system log collectors) for each service set for ms interface under **[edit interfaces *interface-name* services-options]** hierarchy.

Options ***hostname***—Name of the system logging utility host machine. This can be the local Routing Engine or an external server address.

From Junos OS Release 17.4R1, you can configure up to four system log hosts.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *Applying Filters and Services to Interfaces*

host-prefix-only

Syntax

host-prefix-only;

Hierarchy Level

[edit dynamic-profiles interfaces *interface-name* unit *logical-unit-number*],
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information

Statement introduced in Junos OS Release 17.2 on MX Series routers.

Description

(MPC5 and MPC6 cards) Improve datapath performance by allowing only DHCPv4 subscribers that negotiate a 32-bit prefix to come up on the underlying VLAN interface. All DHCP subscribers on the underlying interface must negotiate a 32-bit prefix. Subscribers that negotiate a subnet prefix are not brought up. You can configure this statement for static or dynamic subscribers.

NOTE: You must add or remove this statement before subscribers become active. The configuration fails if you attempt to configure the statement while subscribers are active.

NOTE: You must also configure **demux-source inet** for the logical interface. Only **inet** is supported. A commit error occurs if you specify **demux-source inet6** or **demux-source [inet inet6]**.

Required Privilege Level

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

Related Documentation

- Configuring an IP Demultiplexing Interface on page 316
- Configuring a VLAN Demultiplexing Interface on page 320

 Copyright © 2018, Juniper Networks, Inc.

iccp

Syntax

```
iccp {
  traceoptions; {
    file <filename> <files number> <match regular-expression> <microsecond-stamp>
      <size size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
  }
  local-ip-address ip address;
  session-establishment-hold-time value;
  authentication-key string;
  peer ip-address {
    local-ip-address ip address;
    session-establishment-hold-time value;
    authentication-key string;
    redundancy-group-id-list redundancy-group-id-list;
    liveness-detection;
  }
}
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 10.0.
Support for logical systems introduced in Junos OS Release 14.1.

Description

Configure Interchassis Control Protocol (ICCP) between the multichassis link aggregation group (MC-LAG) peers. ICCP replicates forwarding information, validates configurations, and propagates the operational state of the MC-LAG members.

Default

If you do not include this statement, no ICCP protocol tracing operations are performed.

Options

traceoptions—Set Interchassis Control Protocol (ICCP) tracing options.

local-ip-address—Specify the source address where the ICCP packet is routed.

session-establishment-hold-time—Specify if the chassis takes over as the master at the ICCP session.

authentication-key —Specify TCP Message Digest 5 (MD5) option for an ICCP TCP session.

peer ip-address—Specify the IP address of the peer that hosts an MC-LAG. You must configure ICCP for both peers that host the MC-LAG.

redundancy-group-id-list—Specify the redundancy groups between two ICCP peers.

liveness-detection—Specify Bidirectional Forwarding Detection (BFD) protocol options.

Copyright © 2018, Juniper Networks, Inc.

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

Related Documentation

- Configuring ICCP for MC-LAG on page 214

idle-cycle-flag

Syntax idle-cycle-flag *value*;

Hierarchy Level [edit interfaces e1-*fpc/pic/port*],
[edit interfaces t1-*fpc/pic/port*],
[edit interfaces *interface-name* ds0-options],
[edit interfaces *interface-name* e1-options],
[edit interfaces *interface-name* e3-options],
[edit interfaces *interface-name* serial-options],
[edit interfaces *interface-name* t1-options],
[edit interfaces *interface-name* t3-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers.

Description Configure the value that the DS0, E1, E3, T1, or T3 interface transmits during idle cycles.

NOTE: When configuring E1 or T1 interfaces on 10-port Channelized E1/T1 IQE PICs, the **idle-cycle-flag** statement must be included at the **[edit interfaces e1-*fpc/pic/port*]** or **[edit interfaces t1-*fpc/pic/port*]** hierarchy level as appropriate.

Options *value*—Value to transmit in the idle cycles:

- **flags**—Transmit the value 0x7E.
- **ones**—Transmit the value 0xFF (all ones).

Default: Flags

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

Related Documentation

- *Configuring the E1 Idle Cycle Flag*
- *Configuring the E3 Idle Cycle Flag*
- *Configuring the T1 Idle Cycle Flag*
- *Configuring the T3 Idle Cycle Flag*

 Copyright © 2018, Juniper Networks, Inc.

idle-timeout

Syntax idle-timeout *seconds*;

Hierarchy Level [edit interfaces dl*n* unit *logical-unit-number* dialer-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description On J Series Services Routers with ISDN interfaces, configure the number of seconds the link is idle before losing connectivity.

Options ***seconds***—Time for which the connection can remain idle. For interfaces configured to use a filter for traffic, the idle timeout is based on traffic.

Range: 1 through 429497295

Default: 120 seconds

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

Related Documentation

- *Junos OS Interfaces and Routing Configuration Guide*

Copyright © 2018, Juniper Networks, Inc.

ieee802.1p

Syntax

ieee802.1p premium [*values*];

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile input-priority-map]
[edit interfaces *interface-name* ether-options ethernet-switch-profile ethernet-policer-profile input-priority-map]

Release Information

Statement introduced before Junos Release 7.4.
Statement introduced in Junos OS Release 13.2 for the QFX Series.

Description

For Gigabit Ethernet IQ and 10-Gigabit Ethernet interfaces only, configure premium priority values for IEEE 802.1p input traffic.

Options

values—Define IEEE 802.1p priority values to be treated as premium.

Range: 0 through 7

Required Privilege Level

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

Related Documentation

- *Specifying an Input Priority Map*

 Copyright © 2018, Juniper Networks, Inc.

if-exceeding (Hierarchical Policer)

Syntax

```
if-exceeding {
  bandwidth-limit bps;
  burst-size-limit bytes;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* firewall hierarchical-policer aggregate],
[edit dynamic-profiles *profile-name* firewall hierarchical-policer premium],
[edit firewall hierarchical-policer aggregate],
[edit firewall hierarchical-policer premium]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles ... aggregate]** and **[edit dynamic-profiles ... premium]** hierarchy level introduced in Junos OS Release 11.4.

Description

For M40e, M120, and M320 (with FFPC and SFPC) edge routers and T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs, T4000 routers with Type 5 FPC and Enhanced Scaling Type 4 FPC, specify bandwidth and burst limits for a premium or aggregate component of a hierarchical policer.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

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

Related Documentation

- *Hierarchical Policer Configuration Overview*
- *Hierarchical Policers*
- *aggregate (Hierarchical Policer)*
- bandwidth-limit (Hierarchical Policer) on page 456
- burst-size-limit (Hierarchical Policer) on page 469
- *hierarchical-policer*
- premium (Hierarchical Policer) on page 898

Copyright © 2018, Juniper Networks, Inc.

if-exceeding-pps (Hierarchical Policer)

Syntax

```
if-exceeding-pps {
  pps-limit pps;
  packet-burst packets;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* firewall hierarchical-policer *hierarchical-policer-name* aggregate],
[edit dynamic-profiles *profile-name* firewall hierarchical-policer *hierarchical-policer-name* premium],
[edit firewall hierarchical-policer *hierarchical-policer-name* aggregate],
[edit firewall hierarchical-policer *hierarchical-policer-name* premium]

Release Information

Statement introduced in Junos OS Release 15.2 for MX Series routers with MPCs.

Description

For MX Series routers , **if-exceeding-pps** allows you to configure a packets-per-second (pps)-based trigger for a premium or aggregate component of a hierarchical policer. When applied to the loopback interface (lo0), this kind of trigger can help protect the Routing Engine from DDoS attacks. When applied in other areas, to either transit or control traffic, it is a more fine-grained monitor.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Hierarchical Policer Configuration Overview*
- *Hierarchical Policers*
- *aggregate (Hierarchical Policer)*
- bandwidth-limit (Hierarchical Policer) on page 456
- burst-size-limit (Hierarchical Policer) on page 469
- *hierarchical-policer*
- premium (Hierarchical Policer) on page 898

 Copyright © 2018, Juniper Networks, Inc.

igmp-snooping

List of Syntax

Syntax (EX Series, QFX Series, and NFX Series) on page 655
Syntax (MX Series) on page 655
Syntax (SRX Series) on page 657

Syntax (EX Series, QFX Series, and NFX Series)

```
igmp-snooping {
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable> <match
      regex>;
    flag flag (detail | disable | receive | send);
  }
  vlan (vlan-name | all) {
    data-forwarding {
      source {
        groups group-prefix;
      }
      receiver {
        source-vlans vlan-list;
        install ;
      }
    }
    disable;
    immediate-leave;
    interface interface-name {
      group-limit limit;
      host-only-interface;
      immediate-leave;
      multicast-router-interface;
      static {
        group multicast-ip-address;
      }
    }
    l2-querier {
      source-address ip-address;
      }
    proxy {
      source-address ip-address;
    }
    query-interval seconds;
    query-last-member-interval seconds;
    query-response-interval seconds;
    robust-count number;
    version number;
  }
}
```

Syntax (MX Series)

```
igmp-snooping {
  immediate-leave;
  interface interface-name {
    group-limit limit;
    host-only-interface;
    immediate-leave;
```

Copyright © 2018, Juniper Networks, Inc.

```
      multicast-router-interface;
      static {
        group ip-address {
          source ip-address;
        }
      }
    }
    proxy {
      source-address ip-address;
    }
    query-interval seconds;
    query-last-member-interval seconds;
    query-response-interval seconds;
    robust-count number;
    vlan vlan-id {
      immediate-leave;
      interface interface-name {
        group-limit limit;
        host-only-interface;
        immediate-leave;
        multicast-router-interface;
        static {
          group ip-address {
            source ip-address;
          }
        }
      }
      proxy {
        source-address ip-address;
      }
      query-interval seconds;
      query-last-member-interval seconds;
      query-response-interval seconds;
      robust-count number;
    }
  }
```

 Copyright © 2018, Juniper Networks, Inc.

Syntax (SRX Series)

```
igmp-snooping {
  vlan (all | vlan-name) {
    immediate-leave;
    interface interface-name {
      group-limit range;
      host-only-interface;
      multicast-router-interface;
      immediate-leave;
      static {
        group multicast-ip-address {
        source ip-address;
        }
      }
    }
    l2-querier {
      source-address ip-address;
      }
    proxy {
      source-address ip-address;
    }
    qualified-vlan vlan-id;
    query-interval number;
    query-last-member-interval number;
    query-response-interval number;
  robust-count number;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier>;
    }
  }
}
```

Hierarchy Level

[edit bridge-domains *bridge-domain-name* protocols],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols]
[edit routing-instances *routing-instance-name* protocols]
[edit protocols]

Release Information

Statement introduced in Junos OS Release 8.5.
Statement introduced in Junos OS Release 18.1R1 for SRX1500 devices.
Statement introduced in Junos OS Release 9.1 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for the QFX Series.

Description

Configure IGMP snooping to constrain multicast traffic to only the ports that have receivers attached. IGMP snooping enables the device to selectively send out multicast packets on only the ports that need them. Without IGMP snooping, the device floods the packets on every port. The device listens for the exchange of IGMP messages by the device and the end hosts. In this way, the device builds an IGMP snooping table that has a list of all the ports that have requested a particular multicast group. The factory default configuration enables IGMP snooping on all VLANs.

NOTE: IGMP snooping must be disabled on the device before enabling ISSU.

Default IGMP snooping is disabled on the device.

Options The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation
- *IGMP Snooping in MC-LAG Active-Active Mode*
- *Example: Configuring IGMP Snooping on SRX Series Devices*
- *IGMP Snooping Overview*

ignore

Syntax ignore;

Hierarchy Level [edit interfaces *interface-name* sonet-options trigger *defect*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For ATM over SONET/SDH and SONET/SDH interfaces only, ignore a specific SONET/SDH defect trigger.

Default If you do not include this statement, all defects are honored with no hold time.

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

Related Documentation
- *Configuring SONET/SDH Defect Triggers*
- hold-time (Physical Interface) on page 644

 Copyright © 2018, Juniper Networks, Inc.

ignore-all

Syntax ignore-all;

Hierarchy Level [edit interfaces *interface-name* serial-options dce-options],
[edit interfaces *interface-name* serial-options dte-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Ignore all control leads. You can include the **ignore-all** statement in the configuration only if you do not explicitly enable other signal handling options at the **dte-options** hierarchy level.

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

Related Documentation

- Configuring the Serial Signal Handling on page 344

ignore-l3-incompletes

Syntax ignore-l3-incompletes;

Hierarchy Level [edit interfaces *interface-name* fastether-options],
[edit interfaces *interface-name* gigether-options]

Release Information Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description Ignore the counting of Layer 3 incomplete errors on Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces.

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

Related Documentation

- *Ignoring Layer 3 Incomplete Errors*

Copyright © 2018, Juniper Networks, Inc.

ilmi

Syntax ilmi;

Hierarchy Level [edit interfaces at-*fpc/pic/port* atm-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Enable the router to communicate with directly attached ATM switches and routers. The router uses the VC 0.16 to communicate with the ATM switch or router. Once configured, you can display the IP address and port number of an ATM switch or router using the **show interfaces *interface-name* switch-id** command.

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

Related Documentation

- *Configuring Communication with Directly Attached ATM Switches and Routers*
- *show ilmi*
- *show ilmi statistics*

 Copyright © 2018, Juniper Networks, Inc.

ima-group-options

Syntax

```
ima-group-options {
  differential-delay number;
  frame-length (32 |64 |128 |256);
  frame-synchronization {
    alpha number;
    beta number;
    gamma number;
  }
  minimum-links number;
  symmetry (symmetrical-config-and-operation |
    symmetrical-config-asymmetrical-operation);
  test-procedure {
    ima-test-start;
    ima-test-stop;
    interface name;
    pattern number;
    period number;
  }
  transmit-clock (common |independent);
  version (1.0 |1.1);
}
```

Hierarchy Level [edit interfaces (t1-*fpc/pic/port*:*m*:*n* | e1-*fpc/pic/port*:*n* | t1|e1-*fpc/pic/port*)]

Release Information Statement introduced in Junos OS Release 10.0.
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers.

Description Specify IMA group options.

Options **differential-delay *msec***—Maximum differential delay among links in msec.

Range: 1 through 56

Default: 25

frame-length (32 |64 |128 |256)—IMA frame length in number of cells.

Default: 128

frame-synchronization—IMA group frame synchronization selection.

alpha *number*—Number of consecutive invalid ICP cells for IFSM.

Range: 1 through 2

Default: 2

beta *number*—Number of consecutive errored ICP cells for IFSM.

Range: 1 through 2

Default: 2

gamma *number*—Number of consecutive valid ICP cells for IFSM.

Copyright © 2018, Juniper Networks, Inc.

Range: 1 through 5

Default: 1

minimum-links *number*—IMA group minimum active links.

Range: 1 through 8

Default: 1

symmetry (symmetrical-config-and-operation | symmetrical-config-asymmetrical-operation)—IMA group symmetry mode selection.

test-procedure—Specify an IMA link interface test.

ima-test-start—Start IMA group test.

ima-test-stop—Stop IMA group test.

interface *name*—Interface name of the IMA link to test.

pattern *number*—IMA test pattern.

Range: 1 through 254

Default: 170

period *seconds*—Length of IMA pattern test in seconds.

Range: 1 through 4,294,967,294.

Default: 10

transmit-clock (common |independent)—Transmit clock configuration.

Default: **common**

version (1.0 |1.1)—IMA specification version.

Required Privilege Level

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

Related Documentation

- *ATM Support on Circuit Emulation PICs Overview*
- ima-link-options on page 663
- *Understanding Inverse Multiplexing for ATM*

 Copyright © 2018, Juniper Networks, Inc.

ima-link-options

Syntax
ima-link-options group *g*

Hierarchy Level
[edit interfaces (t1-*fpc/pic/port*:*m*:*n* | e1-*fpc/pic/port*:*n* | t1|e1-*fpc/pic/port*)]

Release Information
Statement introduced in Junos OS Release 10.0.
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers.

Description
Specify an interface as a member of an IMA group.

Options
group *g*—Implying **at-*x/y/g***.

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

Related Documentation

- *ATM Support on Circuit Emulation PICs Overview*
- ima-group-options on page 661
- *Obsolete: Inverse Multiplexing for ATM (IMA) Overview*

inactivity-timeout

Syntax
inactivity-timeout *seconds*;

Hierarchy Level
[edit interfaces *interface-name* services-options]

Release Information
Statement introduced before Junos OS Release 7.4.

Description
For adaptive services interfaces, configure the inactivity timeout period for established flows. The timeout configured in the application protocol definition overrides this value.

Options
seconds—Timeout period, in seconds.
Range: 4 through 86,400 seconds
Default: 30 seconds

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2018, Juniper Networks, Inc.

incoming-map

Syntax

```
incoming-map {
  caller caller-number | accept-all;
}
```

Hierarchy Level

[edit interfaces dl*n* unit *logical-unit-number* dialer-options],
[edit logical-systems *logical-system-name* interfaces dl*n* unit *logical-unit-number* dialer-options]

Release Information

Statement introduced in Junos OS Release 7.5.

Description

On J Series Services Routers with interfaces configured for ISDN, specify the dialer to accept incoming calls.

The remaining statements are explained separately. See CLI Explorer.

NOTE: **The incoming-map statement is mandatory for the router to accept any incoming ISDN calls.**

Required Privilege Level

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

Related Documentation

- *Junos OS Interfaces and Routing Configuration Guide*

Copyright © 2018, Juniper Networks, Inc.

indication

Syntax indication (ignore | normal | require);

Hierarchy Level [edit interfaces *interface-name* serial-options dce-options],
[edit interfaces *interface-name* serial-options dte-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For X.21 interfaces only, configure the from-DCE signal indication.

Options **ignore**—The from-DCE signal is ignored.

normal—Normal indication signal handling as defined by ITU-T Recommendation X.21.

require—The from-DCE signal must be asserted.

Default: normal

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

Related Documentation

- Configuring the Serial Signal Handling on page 344

indication-polarity

Syntax indication-polarity (negative | positive);

Hierarchy Level [edit interfaces *interface-name* serial-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For X.21 interfaces only, configure the indication signal polarity.

Options **positive**—Positive signal polarity.

negative—Negative signal polarity.

Default: positive

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

Related Documentation

- Configuring Serial Signal Polarities on page 347

Copyright © 2018, Juniper Networks, Inc.

ingress-policer-overhead

Syntax ingress-policer-overhead *bytes*;

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number*]

Release Information Statement introduced before Junos OS Release 11.1.
Statement introduced in Junos OS Release 15.1X49-D30 for vSRX.

Description Add the configured number of bytes to the length of a packet entering the interface.

Configure a policer overhead to control the rate of traffic received on an interface. Use this feature to help prevent denial-of-service (DoS) attacks or to enforce traffic rates to conform to the service-level agreement (SLA). When you configure a policer overhead, the configured policer overhead value (bytes) is added to the length of the final Ethernet frame. This calculated length of frame is used to determine the policer or the rate-limiting action.

Traffic policing combines the configured policy bandwidth limits and the burst size to determine how to meter the incoming traffic. If you configure a policer overhead on an interface, Junos OS adds those bytes to the length of incoming Ethernet frames. This added overhead fills each frame closer to the burst size, allowing you to control the rate of traffic received on an interface.

You can configure the policer overhead to rate-limit queues and Layer 2 and Layer 3 policers, for standalone (SA) and high-avalability (HA) deployments. The policer overhead and the shaping overhead can be configured simultaneously on an interface.

NOTE: **vSRX supports policer overhead on Layer 3 policers only.**

The policer overhead applies to all interfaces on the PIC. In the following example, Junos OS adds 10 bytes of overhead to all incoming Ethernet frames on ports ge-0/0/0 through ge-0/0/4.

set chassis fpc 0 pic 0 ingress-policer-overhead 10

NOTE: **vSRX only supports fpc 0 pic 0. When you commit the ingress-policer-overhead statement, the vSRX takes the PIC offline and then back online.**

You need to craft the policer overhead size to match your network traffic. A value that is too low will have minimal impact on traffic bursts. A value that is too high will rate-limit too much of your incoming traffic.

 Copyright © 2018, Juniper Networks, Inc.

In this example, the policer overhead of 255 bytes is configured for ge-0/0/0 through ge-0/0/4. The firewall policer is configured to discard traffic when the burst size is over 1500 bytes. This policer is applied to ge-0/0/0 and ge 0/0/1. Junos OS adds 255 bytes to every Ethernet frame that comes into the configured ports. If, during a burst of traffic, the combined length of incoming frames and the overhead bytes exceeds 1500 bytes, the policer starts to discard further incoming traffic.

```
set chassis fpc 0 pic 0 ingress-policer-overhead 255
set interfaces ge-0/0/0 unit 0 family inet policer input overhead_policer
set interfaces ge-0/0/0 unit 0 family inet address 10.9.1.2/24
set interfaces ge-0/0/1 unit 0 family inet policer input overhead_policer
set interfaces ge-0/0/1 unit 0 family inet address 10.9.2.2/24
set firewall policer overhead_policer if-exceeding bandwidth-limit 32k
set firewall policer overhead_policer if-exceeding burst-size-limit 1500
set firewall policer overhead_policer then discard
```

Options ***bytes***—Number of bytes added to a frame entering an interface.

Range: 0–255 bytes

Default: 0

```
[edit chassis fpc 0 pic 0]
user@host# set ingress-policer-overhead 10;
```

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

Related Documentation

- *ingress-shaping-overhead*
- *Policer Overhead to Account for Rate Shaping Overview*
- *Example: Configuring Policer Overhead to Account for Rate Shaping*
- *Configuring a Policer Overhead*
- *CoS on Enhanced IQ2 PICs Overview*

Copyright © 2018, Juniper Networks, Inc.

ingress-rate-limit

Syntax ingress-rate-limit *rate*;

Hierarchy Level [edit interfaces *interface-name* fastether-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Perform port-based rate limiting on ingress traffic arriving on Fast Ethernet 8-port, 12-port, and 48-port PICs.

Options ***rate***—Traffic rate, in megabits per second (Mbps).

Range: 1 through 100 Mbps

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

Related Documentation

- *Configuring the Ingress Rate Limit*

 Copyright © 2018, Juniper Networks, Inc.

init-command-string

Syntax init-command-string *initialization-command-string*;

Hierarchy Level [edit interfaces umd0 modem-options]

Release Information Statement introduced in Junos OS Release 8.2.

Description For J Series Services Routers, configure the command string used to initialize the USB modem.

When you connect the USB modem to the USB port on a Services Router, the router applies the modem AT commands configured in the **init-command-string** command to the initialization commands on the modem.

For example, the initialization command string **ATS0 = 2\n** configures the USB modem to pick up a call after 2 rings.

If you do not include the **init-command-string** statement, the router applies the default initialization string to the modem.

Options *initialization-command-string*—Specify an initialization command string using the following AT command values:

- **%C0**—Disables data compression.
- **&C1**—Disables reset of the modem when it loses the carrier signal.
- **&Q8**—Enables Microcom Networking Protocol (MNP) error control mode.
- **AT**—Attention. Informs the modem that a command follows.
- **E0**—Disables the display on the local terminal of commands issued to the modem from the local terminal.
- **Q0**—Enables the display of result codes.
- **S0=0**—Disables the auto-answer feature, whereby the modem automatically answers calls.
- **S7=45**—Instructs the modem to wait 45 seconds for a telecommunications service provider (carrier) signal before terminating the call.
- **V1**—Displays result codes as words.

Default: AT S7=45 S0=0 V1 X4 &C1 E0 Q0 &Q8 %C0

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

Related Documentation

- Specifying a USB Modem Interface on J Series Routers on page 350

Copyright © 2018, Juniper Networks, Inc.

initial-route-check

Syntax initial-route-check *seconds*;

Hierarchy Level [edit interfaces dl*n* unit *logical-unit-number* dialer-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description On J Series Services Routers with ISDN interfaces, allows the router to check whether the primary route is up after the initial startup of the router is complete and the timer expires.

Options *seconds*—How long to wait to check if the primary interface is up after the router comes up.

Range: 1 through 300 seconds

Default: 120 seconds

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

Related Documentation

- *ISDN Interfaces Overview*
- *Junos OS Interfaces and Routing Configuration Guide*

 Copyright © 2018, Juniper Networks, Inc.

inner-tag-protocol-id

Syntax

inner-tag-protocol-id *tpid*;

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure the IEEE 802.1Q TPID value to rewrite for the inner tag.

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

On MX Series routers, you can use this statement for Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs.

Default

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

Required Privilege Level

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

Related Documentation

- *Configuring Inner and Outer TPIDs and VLAN IDs*

Copyright © 2018, Juniper Networks, Inc.

inner-vlan-id

Syntax inner-vlan-id *number*;

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

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description For Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers or 100-Gigabit Ethernet Type 5 PIC with CFP, or on Ethernet interfaces on EX Series switches, specify the VLAN ID to rewrite for the inner tag of the final packet.

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

Options *number*—VLAN ID number.

Range: 0 through 4094

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

Related Documentation
- *Configuring Inner and Outer TPIDs and VLAN IDs*

 Copyright © 2018, Juniper Networks, Inc.

inner-vlan-id-range

Syntax inner-vlan-id-range start *start-id* end *end-id*;

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

Release Information Statement introduced in Junos OS Release 9.0.

Description The range of VLAN IDs to be used in the ATM-to-Ethernet interworking cross-connect. Specify the starting VLAN ID and ending VLAN ID.

Options ***start-id***—The lowest VLAN ID to be used.

end-id—The highest VLAN ID to be used.

Range: 32 through 4094

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

Related Documentation

- Configuring ATM-to-Ethernet Interworking on page 282

input

Syntax

```
input {
  service-set service-set-name <service-filter filter-name>;
  post-service-filter filter-name;
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Define one or more input service sets and filters, and one postservice filter to be applied to traffic.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

 Copyright © 2018, Juniper Networks, Inc.

input-list

Syntax input-list [*filter-names*];

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

Release Information Statement introduced in Junos OS Release 7.6.

Description Apply a group of filters to evaluate when packets are received on an interface.

Options **[*filter-names*]**—Name of a filter to evaluate when packets are received on the interface. Up to 16 filters can be included in a filter input list.

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

Related Documentation

- Applying a Filter to an Interface on page 235
- *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*
- *Junos OS Administration Library*
- output-list on page 843

input-policer

Syntax input-policer *policer-name*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* layer2-policer]
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* layer2-policer]

Release Information Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Apply a single-rate two-color policer to the Layer 2 input traffic at the logical interface. The **input-policer** and **input-three-color** statements are mutually exclusive.

Options *policer-name*—Name of the single-rate two-color policer that you define at the **[edit firewall]** hierarchy level.

Usage Guidelines See *Applying Layer 2 Policers to Gigabit Ethernet Interfaces*.

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

Related Documentation

- *Two-Color and Three-Color Policers at Layer 2*
- *Applying Layer 2 Policers to Gigabit Ethernet Interfaces*
- *Configuring a Gigabit Ethernet Policer*
- input-three-color on page 678
- layer2-policer on page 718
- logical-interface-policer on page 747
- output-policer on page 844
- output-three-color on page 846

 Copyright © 2018, Juniper Networks, Inc.

input-priority-map

Syntax

```
input-priority-map {
  ieee802.1p premium [ values ];
}
```

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile]
[edit interfaces *interface-name* ether-options ethernet-switch-profile ethernet-policer-profile]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2 for the QFX Series.

Description

For Gigabit Ethernet IQ and 10-Gigabit Ethernet interfaces only, define the input policer priority map to be applied to incoming frames on this interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Specifying an Input Priority Map*
- output-priority-map on page 845

Copyright © 2018, Juniper Networks, Inc.

input-three-color

Syntax input-three-color *policer-name*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* layer2-policer]
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* layer2-policer]

Release Information Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Apply a single-rate or two-rate three-color policer to the Layer 2 input traffic at the logical interface. The **input-three-color** and **input-policer** statements are mutually exclusive.

Options ***policer-name***—Name of the single-rate or two-rate three-color policer.

Usage Guidelines See *Applying Layer 2 Policers to Gigabit Ethernet Interfaces*.

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

Related Documentation

- *Two-Color and Three-Color Policers at Layer 2*
- *Applying Layer 2 Policers to Gigabit Ethernet Interfaces*
- *Configuring a Gigabit Ethernet Policer*
- input-policer on page 676
- layer2-policer on page 718
- logical-interface-policer on page 747
- output-policer on page 844
- output-three-color on page 846

 Copyright © 2018, Juniper Networks, Inc.

input-vlan-map (Aggregated Ethernet)

Syntax

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

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 8.2.
Starting in Junos OS Release 17.3R1, input-vlan-map for outer vlan is supported for L2 circuit over aggregated Ethernet interfaces for QFX10000 Series switches.

Description

Define the rewrite profile to be applied to incoming frames on this logical interface. On MX Series routers, this statement only applies to aggregated Ethernet inferfaces using Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2 and IQ2-E interfaces and 100-Gigabit Ethernet Type 5 PIC with CFP.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Stacking a VLAN Tag*
- output-vlan-map (Aggregated Ethernet) on page 847

input-vlan-map

Syntax

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

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.
pop-pop, **pop-swap**, **push-push**, **swap-push**, and **swap-swap** statements introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.

Description

For Gigabit Ethernet IQ, 10-Gigabit Ethernet SFPP interfaces, 100-Gigabit Ethernet Type 5 PIC with CFP only as well as Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces, define the rewrite profile to be applied to incoming frames on this logical interface.

The remaining statements are explained separately. See CLI Explorer.

NOTE: Connectivity fault management (CFM) sessions for all interfaces in which **input-vlan-map** is configured are supported only if the interface also has an explicit configuration for **output-vlan-map** as **output-vlan-map pop;**. See output-vlan-map. This configuration is required for all the interfaces in the topology even when the CFM session is on that interface or on a different interface in the data path of the same topology.

Required Privilege Level

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

Related Documentation

- *Stacking a VLAN Tag*
- output-vlan-map on page 848
- *Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support (CLI Procedure)*

 Copyright © 2018, Juniper Networks, Inc.

instance

Syntax instance *vpls-instance-name*;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *name*]

Release Information Statement introduced in Junos OS Release 9.4.

Description Specify the VPLS instance of the default maintenance domain.

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

Related Documentation

- *Configuring Maintenance Intermediate Points (MIPs)*
- maintenance-domain on page 767

interface (Hierarchical CoS Schedulers)

Syntax interface *interface-name*;

Hierarchy Level [edit interfaces interface-set *interface-set-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description Specify an interface that is a member of the interface set. Supported on Ethernet interfaces on an MX Series router, Ethernet interfaces on IQ2E PIC on M Series and T Series routers, and IP demux interfaces on an MX Series router.

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

Related Documentation

- *Class of Service Feature Guide for Routing Devices and EX9200 Switches*

Copyright © 2018, Juniper Networks, Inc.

interface (IEEE 802.1x)

Syntax

```
interface interface-id {
  maximum-requests integer;
  quiet-period seconds;
  reauthentication (disable | interval seconds);
  retries integer;
  server-timeout seconds;
  supplicant (single);
  supplicant-timeout seconds;
  transmit-period seconds;
}
```

Hierarchy Level [edit protocols dot1x authenticator]

Release Information Statement introduced in Junos OS Release 9.3.

Description Use this statement to configure the 802.1x Port-Based Network Access Control protocol-specific Ethernet interface options.

Default The default values are provided for the options below on the respective statement pages.

Options

maximum-requests—Specify the maximum number of retransmission times for an EAPOL Request packet to the client before it times out the authentication session.

quiet-period—Specify the number of seconds the port remains in the wait state following a failed authentication exchange with the client, before reattempting the authentication.

reauthentication—Includes two options:

- **disable**—Periodic reauthentication of the client is disabled.
- **interval**—Specify the periodic reauthentication time interval.

retries—Specify the number of tries after which the port remains in the wait state for **quiet-period** seconds before reattempting the authentication.

server-timeout—Specify the number of seconds the port waits for a reply when relaying a response from the client to the authentication server before timing out and invoking the server-fail action.

supplicant (*single*)—Specify supplicant single mode. See the usage guidelines to configure other modes.

supplicant-timeout—Specify the number of seconds the port waits for a response when relaying a request from the authentication server to the client before resending the request.

 Copyright © 2018, Juniper Networks, Inc.

transmit-period—Specify the number of seconds the port waits before retransmitting the initial EAPOL PDUs to the client.

Required Privilege Level

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

Related Documentation

- *IEEE 802.1x Port-Based Network Access Control Overview*
- authenticator on page 448
- dot1x on page 543

interface (IEEE 802.1ag OAM Connectivity-Fault Management)

Syntax

interface (*interface-name* | ((ge- | xe-) (*fpc/pic/port* | *fpc/pic/port.unit-number* | *fpc/pic/port.unit-number* vlan *vlan-id*)));

Hierarchy Level

[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id*]

Release Information

Statement introduced in Junos OS Release 8.4.

Description

For Ethernet interfaces on M320, MX Series, and T Series routers, configure IEEE 802.1ag Operation, Administration, and Management (OAM) support.

For Gigabit Ethernet interfaces and 10-Gigabit Ethernet interfaces on MX Series routers, configure IEEE 802.1ag Connectivity Fault Management (CFM) support on trunk interface ports.

Starting in Junos OS 17.4R1, you can enable support for IEEE 802.1ag CFM on pseudowire service interfaces by configuring maintenance intermediate points (MIPs) on the pseudowire service interfaces.

NOTE: **The CFM MIP session is supported only on the pseudowire services interface and not on the pseudowire services tunnel interface.**

Options

interface-name—Interface to which the MEP is attached. It could be a physical Ethernet interface, logical Ethernet interface, pseudowire services interfaces, or on a specific VLAN of a trunk port interface (MX Series only).

Required Privilege Level

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

Related Documentation

- *Configuring a MEP to Generate and Respond to CFM Protocol Messages*

Copyright © 2018, Juniper Networks, Inc.

interface (OAM Link-Fault Management)

List of Syntax

Syntax: T, M, MX and ACX Series Routers, SRX Series Firewalls and EX Series Switches on page 684
Syntax: EX Series Switches and NFX Series Devices on page 684

Syntax: T, M, MX and ACX Series Routers, SRX Series Firewalls and EX Series Switches

```
interface interface-name {
  apply-action-profile profile-name;
  link-discovery (active | passive);
  pdu-interval interval;
  pdu-threshold threshold-value;
  remote-loopback;
  event-thresholds {
    frame-error count;
    frame-period count;
    frame-period-summary count;
    symbol-period count;
  }
  negotiation-options {
    allow-remote-loopback;
    no-allow-link-events;
  }
}
```

Syntax: EX Series Switches and NFX Series Devices

```
interface interface-name {
  link-discovery (active | passive);
  pdu-interval interval;
  pdu-threshold threshold-value;
  remote-loopback;
  event-thresholds {
    frame-error count;
    frame-period count;
    frame-period-summary count;
    symbol-period count;
  }
  negotiation-options {
    allow-remote-loopback;
    no-allow-link-events;
  }
}
```

Hierarchy Level

[edit protocols oam ethernet link-fault-management]

Release Information

Statement introduced in Junos OS Release 8.2 for T, M, MX and ACX Series Routers, SRX Series firewalls and EX Series switches.
Statement introduced in Junos OS Release 9.4 for EX Series switches and NFX Series devices.

Description

Configure Ethernet OAM link fault management (LFM) for all interfaces or for specific interfaces.

Copyright © 2018, Juniper Networks, Inc.

For Ethernet interfaces on M320, MX Series, and T Series routers, configure IEEE 802.3ah Operation, Administration, and Management (OAM) support.

Options **interface** ***interface-name***—Interface to be enabled for IEEE 802.3ah link fault management OAM support.

The remaining statements are described separately.

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

Related Documentation

- *Enabling IEEE 802.3ah OAM Support*
- *Example: Configuring Ethernet OAM Link Fault Management*
- *Configuring Ethernet OAM Link Fault Management (CLI Procedure)*

interface (Port Mirroring)

Syntax

```
interface interface-name {
  next-hop address;
}
```

Hierarchy Level [edit forwarding-options port-mirroring family (inet | inet6) output]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the output interface for sending copies of packets elsewhere to be analyzed.

Options ***interface-name***—Name of the interface.

The remaining statements are explained separately. See CLI Explorer.

Usage Guidelines See *Configuring Port Mirroring on M, T MX, and PTX Series Routers*.

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

Copyright © 2018, Juniper Networks, Inc.

interface-down

Syntax interface-down;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management action-profile *profile-name* default-actions]

Release Information Statement introduced in Junos OS Release 8.5.

Description Bring the interface down when a remote MEP connectivity failure is detected.

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

Related Documentation

- *Configuring a CFM Action Profile to Specify CFM Actions for CFM Events*

interface-name

Syntax interface-name;

Hierarchy Level [edit interfaces *interface-name* auto-configure vlan-ranges authentication username-include],
[edit interfaces *interface-name* auto-configure stacked-vlan-ranges authentication username-include],

Release Information Statement introduced in Junos OS Release 10.0.

Description Append the interface name and VLAN ID or stacked VLAN ID to the username string used for authentication. The appended information takes the following format:

- For single VLAN—**<interface-name>:<4-digit-vlan-id>**
- For stack VLANs—**<interface-name>:<4-digit-svlan-id>-<4-digit-vlan-id>**

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

Related Documentation

- *Configuring VLAN Interface Username Information for AAA Authentication*

 Copyright © 2018, Juniper Networks, Inc.

interface-none

Syntax interface-none;

Hierarchy Level [edit protocols protection-group ethernet-ring *ring-name* east-interface]

[edit protocols protection-group ethernet-ring *ring-name* west-interface]

Description Designates port as not used for Ethernet ring protection.

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

Related Documentation

- *Ethernet Ring Protection Switching Overview*
- *Ethernet Ring Protection Using Ring Instances for Load Balancing*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

Copyright © 2018, Juniper Networks, Inc.

interface-range

Syntax

```
interface-range name {
  member-range interface-name-fpc/pic/port to interface-name-fpc/pic/port;
  member interface-name-fpc/pic/port;
  member interface-name-fpc/[low-high]/*;
  member interface-name-fpc/[pic1,pic2,pic3...picN]/port
  /*Common config is added as part of interface-range definition, as follows*/
  mtu 256;
  hold-time up 10;
  ether-options {
    flow-control;
    speed {
      100m;
    }
    802.3ad primary;
  }
}
```

Hierarchy Level [edit interfaces]

Release Information Statement introduced in Junos OS Release 10.0.

Description Specify a set of identical interfaces as an interface group, to which you can apply a common configuration to the entire set of interfaces. This group can consist of both lexical member ranges of interfaces specified using the **member-range *interface-type-fpc/pic/port* to *xx-fpc/pic/port*** option (regex not supported), and of individual or non-sequential members using the **member *interface-type-fpc/pic/port*** option (with regex support to specify the ***fpc/pic/port*** values).

Options **member-range**—Adds interfaces in lexical order. Regex is not supported.

Format:—member-range <start-range> to <end-range>

Example:—member-range ge-0/0/0 to ge-4/0/40;

member—To add individual interfaces or multiple interfaces using regex.

Format:—member <list of interface names>

Example:—member ge-0/0/0;

member ge-0/1/1;

member ge-0/*/*;

member ge-0/[1-10]/0;

member ge-1/[1,3,6,10]/12

 Copyright © 2018, Juniper Networks, Inc.

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

Related Documentation
- Configuring Interface Ranges on page 90

interface-transmit-statistics

Syntax interface-transmit-statistics;

Hierarchy Level [edit interface *interface-name*]

Release Information Statement introduced in Junos OS Release 11.4 R3 for MX Series devices.

Description Configure the interface to report the transmitted load statistics. If this statement is not included in the configuration, the interface statistics show the offered load on the interface, and not the actual transmitted load.

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

Related Documentation
- Improvements to Interface Transmit Statistics Reporting on page 1117
- show interfaces on page 1119

interface-set (Ethernet Interfaces)

Syntax

```
interface-set interface-set-name {
  interface ethernet-interface-name {
    (unit unit-number | vlan-tags-outer vlan-tag);
  }
}
```

Hierarchy Level [edit interfaces]

Release Information Statement introduced in Junos OS Release 8.5.

Description The set of interfaces used to configure hierarchical CoS schedulers on Ethernet interfaces on the MX Series router and IQ2E PIC on M Series and T Series routers.

The remaining statements are described separately.

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

Related Documentation
- *interface-set (Hierarchical Schedulers)*

interface-set (IP Demux Interfaces)

Syntax

```
interface-set interface-set-name {
  interface interface-name {
    unit unit-number;
  }
}
```

Hierarchy Level [edit interfaces]

Release Information Statement introduced in Junos OS Release 9.2.

Description The set of interfaces used to configure hierarchical CoS schedulers for subscribers on IP demux interfaces on the MX Series router.

The remaining statements are described separately.

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

Related Documentation

- *Junos Subscriber Access Configuration Guide*
- *Interfaces Fundamentals for Routing Devices*

 Copyright © 2018, Juniper Networks, Inc.

interface-shared-with

Syntax

interface-shared-with psd*n*;

Hierarchy Level

[edit interfaces ge-*fpc/pic/slot* unit *logical-unit-number*],
[edit interfaces so-*fpc/pic/slot* unit *logical-unit-number*],
[edit interfaces xe-*fpc/pic/slot* unit *logical-unit-number*]

Release Information

Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Assign a logical interface under a shared physical interface to a Protected System Domain (PSD).

Options

n—PSD identification as a numeric value.

Range: **1** through **31**

Required Privilege Level

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

Related Documentation

Copyright © 2018, Juniper Networks, Inc.

interface-status-tlv

Syntax interface-status-tlv [down lower-layer-down];

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management action-profile *profile-name* event]
[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* continuity-check]

Release Information Statement introduced in Junos OS Release 9.6.

Description Defines an **action-profile** consisting of various events and the action. Based on values of **interface-status-tlv** in the received CCM packets, specific action such as *interface-down* can be taken using action-profile options.

Options **down**—When the incoming CCM packet contains interface status TLV with value down, the action will be triggered for this action-profile.

lower-layer-down—When the incoming CCM packet contains interface status TLV with value lower-layer-down, the action will be triggered for this action-profile.

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

Related Documentation

- *Configuring Remote MEP Action Profile Support*

Copyright © 2018, Juniper Networks, Inc.

interface-switch

Syntax

interface-switch *connection-name* {
  interface *interface-name.unit-number*;
}

Hierarchy Level

[edit logical-systems *logical-system-name* protocols connections],
[edit protocols connections]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure Layer 2 switching cross-connects. The cross-connect is bidirectional, so packets received on the first interface are transmitted out the second interface, and those received on the second interface are transmitted out the first.

For Layer 2 switching cross-connects to work, you must also configure MPLS.

Options

connection-name—Connection name (up to 128 characters in Junos 12.3 and later).

interface *interface-name.unit-number*—Interface name. Include the logical portion of the name, which corresponds to the logical unit number.

Required Privilege Level

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

Related Documentation

- *Configuring the CCC Connection for Layer 2 Switching Cross-Connects*
- Defining the Connection for Switching Cross-Connects on page 281
- *MPLS Applications Feature Guide*

Copyright © 2018, Juniper Networks, Inc.

interface-type (Interfaces)

Syntax interface-type (bc | coc1 | ct1 | ct3 | dc | ds |so | t1 | t3);

Hierarchy Level [edit interfaces *interface-range name* no-partition],
[edit interfaces *interface-range name* partition *partition-number*],
[edit interfaces *interface-range name* partition *partition-number* oc-slice *oc-slice-range*],
[edit interfaces *interface-range name* partition *partition-number* timeslot *timeslot-range*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For IQ and IQE interfaces only, configure the sublevel interface type.

Options

bc—Dual—Port Channelized E1 and T1 ISDN PRI interface type. You can specify this interface type at the **[edit interfaces *interface-name* partition *partition-number* timeslot *timeslot-range*]** hierarchy level to create a bearer (B) channel **bc-*pim*/0/*port*:*channel*** interface for each time you want to function as an ISDN PRI B-channel.

coc1—Channelized OC1 interface type. You can specify this interface type at the **[edit interfaces *interface-name* partition *partition-number* oc-slice *oc-slice-range* interface-type coc12-*fpc/pic/port*]** hierarchy level.

ct1—Channelized T1 interface type. You can specify this interface type at the **[edit interfaces *interface-name* partition *partition-number* interface-type ct3-*fpc/pic/port*<:*channel*>]** hierarchy level.

ct3—Channelized T3 interface type. You can specify this interface type at the **[edit interfaces *interface-name* partition *partition-number* oc-slice *oc-slice-range* interface-type coc1-*fpc/pic/port*:*channel* no-partition]** hierarchy level.

dc—Dual-Port Channelized E1 and T1 ISDN PRI interface type. You can specify this interface type at the **[edit interfaces *interface-name* partition *partition-number* timeslot *timeslot-range*]** hierarchy level to create a (D) channel **dc-*pim*/0/*port*** to control the B-channels.

ds—DS0 interface type. You can specify this interface type at the **[edit interfaces *interface-name*partition *partition-number* interface-type (ce1-*fpc/pic/port* | ct1-*fpc/pic/port*<:*channel*>)]** hierarchy level.

so—SONET/SDH interface type. You can specify this interface type at the **[edit interfaces *interface-name* partition *partition-number* oc-slice *oc-slice-range* interface-type coc12-*fpc/pic/port*]** hierarchy level.

t1—T1 interface type. You can specify this interface type at the **[edit interfaces *interface-name* partition *partition-number* oc-slice *oc-slice-range* interface-type (coc12-*fpc/pic/port* | coc1-*fpc/pic/port*)]** hierarchy level.

 Copyright © 2018, Juniper Networks, Inc.

t3—T3 interface type. You can specify this interface type at the **[edit interfaces *interface-name* partition *partition-number* oc-slice*oc-slice-range* interface-type (coc12-*fpc/pic/port* | coc1-*fpc/pic/port*:*channel* no-partition)]** hierarchy level.

Required Privilege Level

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

Related Documentation

- *Channelized E1 IQ and IQE Interfaces Overview*
- *Channelized OC12/STM4 IQ and IQE Interfaces Overview*
- *Configuring Channelized T3 IQ Interfaces*

interfaces

List of Syntax

Syntax (QFX Series) on page 696
Syntax (EX Series, MX Series and T Series) on page 696

Syntax (QFX Series)

```
interfaces interface-name {
  no-mac-learning;
}
```

Syntax (EX Series, MX Series and T Series)

```
interfaces { ... }
```

QFX Series

[edit ethernet-switching-options]

EX Series, MX Series and T Series

[edit]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description

Configure settings for interfaces that have been assigned to family **ethernet-switching**.

Default

The management and internal Ethernet interfaces are automatically configured. You must configure all other interfaces.

Options

interface-name —Name of an interface that is configured for **family ethernet-switching**.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- Physical Interface Configuration Statements Overview on page 64
- *Configuring Aggregated Ethernet Link Protection*

Copyright © 2018, Juniper Networks, Inc.

interfaces (Static and Dynamic Subscribers)

Syntax

```
interfaces {
  interface-name {
    unit logical-unit-number {
      actual-transit-statistics;
      auto-configure {
        agent-circuit-identifier {
          dynamic-profile profile-name;
        }
        line-identity {
          include {
            accept-no-ids;
            circuit-id;
            remote-id;
          }
          dynamic-profile profile-name;
        }
      }
      family family {
        access-concentrator name;
        address address;
        direct-connect;
        duplicate-protection;
        dynamic-profile profile-name;
        filter {
          adf {
            counter;
            input-precedence precedence;
            not-mandatory;
            output-precedence precedence;
            rule rule-value;
          }
          input filter-name {
            precedence precedence;
            shared-name  filter-shared-name;
          }
          output filter-name {
            precedence precedence;
            shared-name  filter-shared-name;
          }
        }
        max-sessions number;
        max-sessions-vsa-ignore;
        rpf-check {
          mode loose;
        }
        service {
          input {
            service-set service-set-name {
              service-filter filter-name;
            }
            post-service-filter filter-name;
          }
```

```
        output {
          service-set service-set-name {
            service-filter filter-name;
          }
        }
      }
      service-name-table table-name
      short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
        maximum-seconds>;
      unnumbered-address interface-name <preferred-source-address address>;
    }
    filter {
      input filter-name (
        precedence precedence;
        shared-name  filter-shared-name;
      }
      output filter-name {
        precedence precedence;
        shared-name  filter-shared-name;
      }
    }
    host-prefix-only;
    ppp-options {
      chap;
      pap;
    }
    proxy-arp;
    service {
      pcef  pcef-profile-name {
        activate rule-name | activate-all;
      }
    }
    vlan-id;
    vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
  }
  vlan-tagging;
}
interface-set interface-set-name {
  interface interface-name {
    unit logical unit number {
      advisory-options {
        downstream-rate rate;
        upstream-rate rate;
      }
    }
  }
  pppoe-underlying-options {
    max-sessions number;
  }
}
demux0 {
  unit logical-unit-number {
    demux-options {
      underlying-interface interface-name
    }
    family family {
```


Copyright © 2018, Juniper Networks, Inc.

```
      access-concentrator name;
      address address;
      direct-connect;
      duplicate-protection;
      dynamic-profile profile-name;
      demux-source {
        source-prefix;
      }
      filter {
        input filter-name (
          precedence precedence;
          shared-name filter-shared-name;
        }
        output filter-name {
          precedence precedence;
          shared-name filter-shared-name;
        }
      }
      mac-validate (loose | strict):
      max-sessions number;
      max-sessions-vsa-ignore;
      rpf-check {
        fail-filter filter-name;
        mode loose;
      }
      service-name-table table-name
      short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
        maximum-seconds>;
      unnumbered-address interface-name <preferred-source-address address>;
    }
    filter {
      input filter-name;
      output filter-name;
    }
    vlan-id number;
    vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
  }
}
pp0 {
  unit logical-unit-number {
    keepalives interval seconds;
    no-keepalives;
    pppoe-options {
      underlying-interface interface-name;
      server;
    }
    ppp-options {
      aaa-options aaa-options-name;
      authentication [ authentication-protocols ];
      chap {
        challenge-length minimum minimum-length maximum maximum-length;
        local-name name;
      }
      ignore-magic-number-mismatch;
      initiate-ncp (dual-stack-passive | ipv6 | ip)
      ipcp-suggest-dns-option;
```

Copyright © 2018, Juniper Networks, Inc.

```
          mru size;
          mtu (size | use-lower-layer);
          on-demand-ip-address;
          pap;
          peer-ip-address-optional;
          local-authentication {
            password password;
            username-include {
              circuit-id;
              delimiter character;
              domain-name name;
              mac-address;
              remote-id;
            }
          }
        }
        family inet {
          unnumbered-address interface-name;
          address address;
          service {
            input {
              service-set service-set-name {
                service-filter filter-name;
              }
              post-service-filter filter-name;
            }
            output {
              service-set service-set-name {
                service-filter filter-name;
              }
            }
          }
          filter {
            input filter-name {
              precedence precedence;
              shared-name filter-shared-name;
            }
            output filter-name {
              precedence precedence;
              shared-name filter-shared-name;
            }
          }
        }
      }
    }
  }
```

Hierarchy Level [edit dynamic-profiles *profile-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Define interfaces for dynamic profiles.

 Copyright © 2018, Juniper Networks, Inc.

Options *interface-name*—The interface variable (**$junos-interface-ifd-name**). The interface variable is dynamically replaced with the interface the DHCP client accesses when connecting to the router.

NOTE: **Though we do not recommend it, you can also enter the specific name of the interface you want to assign to the dynamic profile.**

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

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

Related Documentation

- *Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles*
- *Configuring Dynamic PPPoE Subscriber Interfaces*
- *Configuring Dynamic VLANs Based on Agent Circuit Identifier Information*
- *DHCP Subscriber Interface Overview*
- *Subscribers over Static Interfaces Configuration Overview*
- Demultiplexing Interface Overview on page 313

interleave-fragments

Syntax interleave-fragments;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services interfaces only, interleave long packets with high-priority packets.

Allows small delay-sensitive packets, such as Voice over IP (VoIP) packets, to interleave with long fragmented packets. This minimizes the latency of delay-sensitive packets.

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2018, Juniper Networks, Inc.

interval

Syntax interval (100ms | 10m | 10ms | 10s | 1m | 1s);

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* continuity-check]

Release Information Statement introduced in Junos OS Release 8.4.
Option **10ms** introduced in Junos OS Release 9.1.
Third-party interoperability during a unified in-service software upgrade (ISSU) introduced in Junos OS Release 17.1.

Description Configure the interval between successive transmissions of continuity check messages (CCMs) as part of the connectivity fault detection strategy. When the receiving maintenance association end point (MEP) does not receive a CCM at the configured interval, the **loss-threshold** statement determines how many CCMs can be lost before the sending MEP is marked as down. The **hold-interval** statement then determines the frequency at which the database of MEPs in the maintenance association (MA) is flushed in the absence of updates.

During a unified in-service software upgrade (ISSU), Junos OS connectivity fault management (CFM) works when the peer device is not a Juniper Networks router. Interoperating with the router of another vendor, the Juniper Networks router retains session information and continues to transmit CCM (continuity check message) PDUs during the unified ISSU upgrade. For this feature to work, you must enable Packet Forwarding Engine keepalives with the **hardware-assisted-keepalives** statement, and configure the interval between CCMs to be 1 second with **interval** statement.

NOTE: For the continuity check message interval to be configured for 10 milliseconds, periodic packet management (PPM) runs on the Routing Engine and Packet Forwarding Engine by default. You can disable PPM only on the Packet Forwarding Engine. To disable PPM on the Packet Forwarding Engine, use the **no-delegate-processing** statement at the **[edit routing-options ppm]** hierarchy level.

NOTE: A continuity check interval of 10 milliseconds is not supported for CFM sessions over a label-switched interface (LSI).

Options **100ms**—100 milliseconds.

10m—10 minutes.

10ms—10 milliseconds.

Copyright © 2018, Juniper Networks, Inc.

10s—10 seconds.

1m—1 minute.

1s—1 second.

Default: 1m

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

Related Documentation

- *Continuity Check Protocol Parameters Overview*
- *Configuring Continuity Check Protocol Parameters for Fault Detection*
- *Configuring Connectivity Fault Management for Interoperability During Unified In-Service Software Upgrades*

inverse-arp

Syntax
inverse-arp;

Hierarchy Level
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* multipoint-destination *destination*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address* multipoint-destination *destination*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description
For ATM encapsulation, enable responses to receive inverse ATM ARP requests. For Frame Relay encapsulation, enable responses to receive inverse Frame Relay ARP requests.

Default
Inverse ARP is disabled on all ATM and Frame Relay interfaces.

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

Related Documentation

- *Configuring Inverse ATM1 or ATM2 ARP*
- *Configuring Inverse Frame Relay ARP*

Copyright © 2018, Juniper Networks, Inc.

invert-data

Syntax invert-data;

Hierarchy Level [edit interfaces e1-*fpc/pic/port*],
[edit interfaces t1-*fpc/pic/port*],
[edit interfaces *interface-name* ds0-options],
[edit interfaces *interface-name* e1-options],
[edit interfaces *interface-name* t1-options],
[edit interfaces *interface-name* e3-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers.

Description Invert the transmission of unused data bits on the DS0, E1, E3, and T1 interface.

NOTE: When configuring E1 or T1 interfaces on 10-port Channelized E1/T1 IQE PICs, the **invert-data** statement must be included at the **[edit interfaces e1-*fpc/pic/port*]** or **[edit interfaces t1-*fpc/pic/port*]** hierarchy level as appropriate.

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

Related Documentation

- *Configuring E1 Data Inversion*
- *Configuring E3 Data Inversion*
- *Configuring T1 Data Inversion*

 Copyright © 2018, Juniper Networks, Inc.

ipsec-sa

Syntax

ipsec-sa *sa-name*;

Hierarchy Level

[edit interfaces es-*fpc/pic/port* unit *logical-unit-number* family inet],
[edit logical-systems *logical-system-name* interfaces es-*fpc/pic/port* unit *logical-unit-number* family inet]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the IP Security (IPsec) security association (SA) name associated with the interface.

Options

sa-name—IPsec security association name.

Required Privilege Level

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*
- *Junos OS Administration Library*

Copyright © 2018, Juniper Networks, Inc.

isdn-options

Syntax

```
isdn-options {
   bchannel-allocation (ascending | descending);
   calling-number number;
   incoming-called-number number <reject>;
   spid1 spid-string;
   spid2 spid-string;
   static-tei-val value;
   switch-type (att5e | etsi | ni1 | ntdms100 | ntt);
   t310 seconds;
   tei-option (first-call | power-up);
}
```

Hierarchy Level

[edit interfaces br-*pim*/0/*port*],
[edit interfaces ct1-*pim*/0/*port*],
[edit interfaces ce1-*pim*/0/*port*]

Release Information

Statement introduced before Junos OS Release 7.4.
bchannel-allocation option added in Junos OS Release 8.3.

Description

For J Series Services Routers only. Specify the ISDN options for configuring ISDN interfaces for group and user sessions.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Configuring ISDN Physical Interface Properties*
- *Allocating B-Channels for Dialout*
- *Junos OS Interfaces and Routing Configuration Guide*

Copyright © 2018, Juniper Networks, Inc.

iteration-count

Syntax iteration-count *count-value*;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id* remote-mep *remote-mep-id* sla-iterator-profile *profile-name*]

Release Information Statement introduced in Junos OS Release 11.1.

Description Configure the number of iterations for which the connection partakes in the iterator for acquiring SLA measurements.

Options *count-value*—Number of iterations for which the connection should partake in the iterator for acquiring SLA measurements.

Range: 1 through 65,535

Default: 0 (or infinite iterations)

Required Privilege Level Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- sla-iterator-profile on page 969
- *Configuring a Remote MEP with an Iterator Profile*

iteration-period

Syntax iteration-period *iteration-period-value*;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles *profile-name*]

Release Information Statement introduced in Junos OS Release 11.1.
Statement introduced in Junos OS Release 11.4 for EX Series switches.

Description Configure the iteration period, which is the maximum number of cycles per iteration (that is, the number of connections registered to an iterator cannot exceed this value).

Options *iteration-period-value*—Maximum number of cycles per iteration.

Range: 1 through 2000

Default: 2000

Required Privilege Level Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- *Configuring an Iterator Profile*
- *Configuring an Iterator Profile on a Switch (CLI Procedure)*

 Copyright © 2018, Juniper Networks, Inc.

keep-address-and-control

Syntax keep-address-and-control;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description For interfaces with encapsulation type PPP CCC, do not remove the address and control bytes before encapsulating the packet into a tunnel.

Default If you do not include this statement, address and control bytes are removed before encapsulating the packet into a tunnel.

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

Related Documentation

- Disabling the Removal of Address and Control Bytes on page 225

keepalives

Syntax keepalives <interval *seconds*> <down-count *number*> <up-count *number*>;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Enable the sending of keepalives on a physical interface configured with PPP, Frame Relay, or Cisco HDLC encapsulation.

For ATM2 IQ interfaces only, you can enable keepalives on a logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:

- **atm-ppp-llc**—PPP over AAL5 LLC encapsulation.
- **atm-ppp-vc-mux**—PPP over AAL5 multiplex encapsulation.

Default Sending of keepalives is enabled by default. The default keepalive interval is 10 seconds for PPP, Frame Relay, or Cisco HDLC. The default down-count is 3 and the default up-count is 1 for PPP or Cisco HDLC.

Options **down-count *number***—The number of keepalive packets a destination must fail to receive before the network takes down a link.

Range: 1 through 255

Default: 3

interval *seconds*—The time in seconds between successive keepalive requests.

Range: 1 through 32767 seconds

Default: 10 seconds

up-count *number*—The number of keepalive packets a destination must receive to change a link's status from down to up.

Range: 1 through 255

Default: 1

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

Related Documentation

- Configuring Keepalives on page 134
- *Configuring Frame Relay Keepalives*
- *Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface*

 Copyright © 2018, Juniper Networks, Inc.

key

Syntax key *number*;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description For Adaptive Services PICs on M Series routers (except the M320 and M120 routers), identify an individual traffic flow within a tunnel, as defined in RFC 2890, *Key and Sequence Number Extensions to GRE*.

Options *number*—Value of the key.

Range: 0 through 4,294,967,295

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2018, Juniper Networks, Inc.

l2tp-interface-id

Syntax

l2tp-interface-id *name*;
(dedicated | shared);

Hierarchy Level

[edit interfaces sp-*fpc/pic/port* unit *logical-unit-number* interface],
[edit logical-systems *logical-system-name* interfaces sp-*fpc/pic/port* unit *logical-unit-number* interface]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the L2TP options for configuring logical interfaces for group and user sessions.

Options

(**dedicated** | **shared**)—Specifies whether a logical interface can host one (dedicated) or multiple (shared) sessions at one time.

name—Interface identifier that must be replicated at the **[edit access profile *name*]** hierarchy level.

Required Privilege Level

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2018, Juniper Networks, Inc.

lacp (Protocols)

Syntax

```
lacp {
  traceoptions {
    file <filename> <files number> <size size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
  }
  fast-hello-issu;
  ppm (Ethernet Switching) centralized;
}
```

Hierarchy Level

[edit protocols]

Release Information

Statement introduced in Junos OS Release 9.3.
The **ppm centralized** option introduced in Junos OS Release 9.4.
The **fast-hello-issu** option introduced in Junos OS Release 14.1.

Description

On MX and T Series routers, you can specify periodic packet management (PPM) as centralized. By default, the PPM is distributed.

MX Series routers support Link Aggregation Control Protocol (LACP) with fast hellos during unified ISSU. This support is disabled by default. You must enable the **fast-hello-issu** option on the main router and on the peer routers before starting unified ISSU. Note that the peer router must also be an MX Series router for this functionality to work.

Default

Distributed PPM processing is enabled for all packets that use PPM.

Options

ppm—Set PPM to centralized.

fast-hello-issu—Enable LACP with fast hellos during unified ISSU.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Tracing LACP Operations*

Copyright © 2018, Juniper Networks, Inc.

lacp (802.3ad)

Syntax

```
lacp {
  port-priority port-priority;
}
```

Hierarchy Level

[edit interfaces *interface-name* fastether-options 802.3ad],
[edit interfaces *interface-name* gigether-options 802.3ad]

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Configure the Link Aggregation Control Protocol (LACP) port priority for Ethernet interfaces.

Options

port-priority—Priority for being elected as the active port to collect and distribute traffic. A smaller value indicates a higher priority for selection.

Range: 0 through 65,535

Default: 127

Required Privilege Level

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

Related Documentation

- *Configuring LACP for Aggregated Ethernet Interfaces*
- port-priority on page 889

Copyright © 2018, Juniper Networks, Inc.

lacp (Aggregated Ethernet)

List of Syntax

Syntax (NFX Series) on page 715
Syntax (EX Series) on page 715

Syntax (NFX Series)

```
lacp (active | passive) {
  admin-key key;
  fast-failover;
  link-protection {
    disable;
    (revertive | non-revertive);
  }
  periodic interval
  system-ID mac-address;
  system-priority priority;
  force-up;
}
```

Syntax (EX Series)

```
lacp {
  (active | passive);
  admin-key key;
  accept-data;
  fast-failover;
    link-protection {
    disable;
    (revertive |non-revertive);
  }
  periodic interval;
  system-id mac-address;
  system-priority priority;
}
```

Hierarchy Level (EX Series)

[edit interfaces ae*x* aggregated-ether-options]
[edit logical-systems *logical-system-name* interfaces ae*X* aggregated-ether-options]

Hierarchy Level (NFX Series)

[edit interfaces *interface-name*aggregated-ether-options]

Release Information

Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure the Link Aggregation Control Protocol (LACP) parameters for interfaces. The remaining statement is explained separately.

For EX Series, when you configure the **accept-data** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level, the router processes packets received on a member link irrespective of the LACP state if the aggregated Ethernet bundle is up.

Copyright © 2018, Juniper Networks, Inc.

NOTE: When you configure the **accept-data** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level, this behavior occurs:

- By default, the **accept-data** statement is not configured when LACP is enabled.
- You can configure the **accept-data** statement to improve convergence and reduce the number of dropped packets when member links in the bundle are enabled or disabled.
- When LACP is down and a member link receives packets, the router or switch does not process packets as defined in the IEEE 802.1ax standard. According to this standard, the packets should be dropped, but they are processed instead because the **accept-data** statement is configured.

NOTE: The force-up statement is not supported on QFX10002 switches.

Default If you do not specify LACP as either **active** or **passive**, LACP remains passive.

Options **active**—Initiate transmission of LACP packets.

admin-key ***number***—Specify an administrative key for the router or switch.

NOTE: You must also configure multichassis link aggregation (MC-LAG) when you configure the **admin-key**.

fast-failover—Specify to override the IEEE 802.3ad standard and allow the standby link to receive traffic. Overriding the default behavior facilitates subsecond failover.

passive—Respond to LACP packets.

The remaining statements are explained separately. See CLI Explorer.

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

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- *Configuring Link Aggregation*
- *Configuring Aggregated Ethernet LACP (CLI Procedure)*
- *Understanding Aggregated Ethernet Interfaces and LACP for Switches*
- *Configuring LACP for Aggregated Ethernet Interfaces*

Copyright © 2018, Juniper Networks, Inc.

layer2-policer

Syntax

```
layer2-policer {
  input-policer policer-name;
  input-three-color policer-name;
  output-policer policer-name;
  output-three-color policer-name;
}
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

For 1-Gigabit Ethernet and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces on M Series, MX Series, and T Series routers, and for aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces on EX Series switches, apply Layer 2 logical interface policers. The following policers are supported:

- Two-color
- Single-rate tricolor marking (srTCM)
- Two-rate tricolor marking (trTCM)

Two-color and tricolor policers are configured at the **[edit firewall]** hierarchy level.

Options

input-policer *policer-name*—Two-color input policer to associate with the interface. This statement is mutually exclusive with the **input-three-color** statement.

input-three-color *policer-name*—Tricolor input policer to associate with the interface. This statement is mutually exclusive with the **input-policer** statement.

output-policer *policer-name*—Two-color output policer to associate with the interface. This statement is mutually exclusive with the **output-three-color** statement.

output-three-color *policer-name*—Tricolor output policer to associate with the interface. This statement is mutually exclusive with the **output-policer** statement.

Required Privilege Level

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

Related Documentation

- *Applying Layer 2 Policers to Gigabit Ethernet Interfaces*
- *Configuring Gigabit Ethernet Two-Color and Tricolor Policers*

Copyright © 2018, Juniper Networks, Inc.

lcp-max-conf-req

Syntax lcp-max-conf-req *number*

Hierarchy Level [edit interfaces so-*fpc/pic/port* unit *number* ppp-options]

Release Information Statement introduced in Junos OS Release 9.6.

Description Set the maximum number of LCP Configure-Requests to be sent, after which the router goes to LCP down state.

Options ***number***—From 0 to 65,535, where 0 means send infinite LCP Configure-Requests, and any other value specifies the maximum number LCP Configure-Requests to send and then stop sending.

Default—254

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

Related Documentation

- Configuring the LCP Configure-Request Maximum Sent on page 195
- **ppp-options on page 891**

lcp-restart-timer

Syntax lcp-restart-timer *milliseconds*;

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

Release Information Statement introduced in Junos OS Release 8.1.

Description For interfaces with PPP, PPP TCC, PPP over Ethernet, PPP over ATM, and PPP over Frame Relay encapsulations, configure a restart timer for the Link Control Protocol (LCP) component of a PPP session.

Options ***milliseconds***—The time, in milliseconds, between successive LCP configuration requests.

Range: 20 through 10000 milliseconds

Default: 3 seconds

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

Related Documentation

- Configuring the PPP Restart Timers on page 193

 Copyright © 2018, Juniper Networks, Inc.

level

Syntax level *number*;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name*]

Release Information Statement introduced in Junos OS Release 8.4.
Statement introduced in junos os release 12.1X48 for PTX Series Packet Transport Routers.

Description A number used in connectivity fault management (CFM) messages to identify the maintenance association. The number is embedded in each of the CFM frames. CFM messages within a given level are processed by maintenance end points (MEPs) at the same level. For example, the operator domain can be level 0, the provider domain can be level 3, and the customer domain can be level 7.

Options *number*—A number used to identify the maintenance domain to which the CFM message belongs.

Range: 0 through 7

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

Related Documentation
- *Creating a Maintenance Domain*

Copyright © 2018, Juniper Networks, Inc.

line-encoding

Syntax line-encoding (ami | b8zs);

Hierarchy Level [edit interfaces ct1-*fpc/pic/port*],
[edit interfaces *interface-name* t1-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers.

Description Set the line encoding format on the T1 interface.

NOTE: When configuring CT1 interfaces on the 10-port Channelized E1/T1 IQE PIC, the **line-encoding** statement must be included at the **[edit interfaces ct1-*fpc/pic/port*]** hierarchy level.

Default The default line encoding is B8ZS.

Options **ami**—Use Alternate Mark Inversion (AMI) line encoding.

b8zs—Use bipolar with 8-zeros substitution (B8ZS) line encoding.

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

Related Documentation

- *Configuring T1 Line Encoding*

 Copyright © 2018, Juniper Networks, Inc.

line-protocol

Syntax line-protocol *protocol*;

Hierarchy Level [edit interfaces *interface-name* serial-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For serial interfaces only, configure the line protocol.

Options *protocol*—You can specify the one of the following line protocols:

- **eia530**—Line protocol EIA-530
- **v.35**—Line protocol V.35
- **x.21**—Line protocol X.21

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

Related Documentation

- Configuring the Serial Line Protocol on page 338

Copyright © 2018, Juniper Networks, Inc.

line-rate

Syntax line-rate *line-rate*;

Hierarchy Level [edit interfaces *interface-name* shdsl-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* shdsl-options]

Release Information Statement introduced in Junos OS Release 7.4.

Description For J Series Services Routers only, configure the SHDSL line rate.

Options ***line-rate***—SHDSL line rate, in Kbps. Possible values are:

2-wire (Kbps): **192**, **256**, **320**, **384**, **448**, **512**, **576**, **640**, **704**, **768**, **832**, **896**, **960**, **1024**, **1088**, **1152**, **1216**, **1280**, **1344**, **1408**, **1472**, **1536**, **1600**, **1664**, **1728**, **1792**, **1856**, **1920**, **1984**, **2048**, **2112**, **2176**, **2240**, **2304**, **auto**

4-wire (Kbps): **384**, **512**, **640**, **768**, **896**, **1024**, **1152**, **1280**, **1408**, **1536**, **1664**, **1792**, **1920**, **2048**, **2176**, **2304**, **2432**, **2560**, **2688**, **2816**, **2944**, **3072**, **3200**, **3328**, **3456**, **3584**, **3712**, **3840**, **3968**, **4096**, **4224**, **4352**, **4480**, **4608**

Default: For 2-wire mode, **auto**; for 4-wire mode, **4608** Kbps

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

 Copyright © 2018, Juniper Networks, Inc.

linear-red-profile

Syntax linear-red-profile *profile-name*;

Hierarchy Level [edit interfaces at-*fpc/pic/port* atm-options scheduler-maps *map-name* forwarding-class *class-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For ATM2 IQ interfaces only, assign a linear RED profile to a specified forwarding class. To define the linear RED profiles, include the linear-red-profiles statement at the **[edit interfaces at-*fpc/pic/port* atm-options]** hierarchy level.

Default If you do not include either the **epd-threshold** or the **linear-red-profile** statement in the forwarding class configuration, the Junos OS uses an EPD threshold based on the available bandwidth and other parameters.

Options ***profile-name***—Name of the linear RED profile.

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

Related Documentation

- *Configuring an ATM Scheduler Map*
- linear-red-profiles on page 726
- *Applying Scheduler Maps to ATM Interfaces*
- *epd-threshold*

linear-red-profiles

Syntax

```
linear-red-profiles profile-name {
  high-plp-threshold percent;
  low-plp-threshold percent;
  queue-depth cells;
}
```

Hierarchy Level [edit interfaces at-*fpc/pic/port* atm-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For ATM2 IQ interfaces only, define CoS virtual circuit drop profiles for RED. When a packet arrives, RED checks the queue fill level. If the fill level corresponds to a nonzero drop probability, the RED algorithm determines whether to drop the arriving packet.

Options ***profile-name***—Name of the drop profile.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *ATM2 IQ VC Tunnel CoS Components Overview*
- *Configuring Linear RED Profiles on ATM Interfaces*

link-adjacency-loss

Syntax link-adjacency-loss;

Hierarchy Level [edit protocols oam ethernet link-fault-management action-profile event]

Release Information Statement introduced in Junos OS Release 8.5.

Description Loss of adjacency with IEEE 802.3ah link-fault management peer event. When included, the loss-of-adjacency event triggers the action specified under the **action** statement.

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

Related Documentation

- *Monitoring the Loss of Link Adjacency*

 Copyright © 2018, Juniper Networks, Inc.

link-discovery

Syntax link-discovery (active | passive);

Hierarchy Level [edit protocols oam ethernet link-fault-management interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.2.

Description For Ethernet interfaces on EX Series switches, and M320, M120, MX Series, and T Series routers, specify the discovery mode used for IEEE 802.3ah Operation, Administration, and Management (OAM) support. The discovery process is triggered automatically when OAM 802.3ah functionality is enabled on a port. Link monitoring is done when the interface sends periodic OAM PDUs.

Options **(active | passive)**—Passive or active mode. In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality. In passive mode, the peer initiates the discovery process. Once the discovery process is initiated, both sides participate in discovery.

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

Related Documentation
- *Configuring Link Discovery*

link-down

Syntax link-down;

Hierarchy Level [edit protocols oam ethernet link-fault-management]

Release Information Statement introduced in Junos OS Release 8.5.

Description Mark the interface down for transit traffic.

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

Related Documentation
- *Specifying the Actions to Be Taken for Link-Fault Management Events*

Copyright © 2018, Juniper Networks, Inc.

link-event-rate

Syntax

```
link-event-rate {
  frame-error count;
  frame-period count;
  frame-period-summary count;
  symbol-period count;
}
```

Hierarchy Level [edit protocols oam ethernet link-fault-management action-profile event]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure the number of link-fault management events per second.

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

Related Documentation

- *Configuring Threshold Values for Fault Events in an Action Profile*

 Copyright © 2018, Juniper Networks, Inc.

link-fault-management

Syntax

```
link-fault-management {
  action-profile profile-name {
    action {
      link-down;
      send-critical-event;
      syslog;
    }
    event {
      link-adjacency-loss;
      link-event-rate {
        frame-error count;
        frame-period count;
        frame-period-summary count;
        symbol-period count;
      }
      protocol-down;
    }
  }
  interface interface-name {
    apply-action-profile profile-name;
    link-discovery (active | passive);
    loopback-tracking;
    pdu-interval interval;
    pdu-threshold threshold-value;
    remote-loopback;
    event-thresholds {
      frame-error count;
      frame-period count;
      frame-period-summary count;
      symbol-period count;
    }
    negotiation-options {
      allow-remote-loopback;
      no-allow-link-events;
    }
  }
}
```

Hierarchy Level [edit protocols oam ethernet]

Release Information Statement introduced in Junos OS Release 8.2.

Description For Ethernet interfaces on M320, M120, MX Series, and T Series routers and EX Series switches, specify fault signaling and detection for IEEE 802.3ah Operation, Administration, and Management (OAM) support.

The remaining statements are explained separately. See CLI Explorer.

Copyright © 2018, Juniper Networks, Inc.

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

Related Documentation
- *Enabling IEEE 802.3ah OAM Support*

link-layer-overhead

Syntax link-layer-overhead *percent*;

Hierarchy Level [edit interfaces *interface-name* mlfr-uni-nni-bundle-options],
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For AS PIC or MultiServices PIC link services IQ interfaces (**lsq**) only, configure the percentage of total bundle bandwidth to be set aside for link-layer overhead.

Options ***percent***—Percentage of total bundle bandwidth to be set aside for link-layer overhead.

Range: 0 through 50 percent

Default: 4 percent

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

Related Documentation
- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2018, Juniper Networks, Inc.

link-mode

Syntax link-mode *mode* (automatic | full-duplex | half-duplex);

Hierarchy Level [edit interfaces *interface-name*],
[edit interfaces *interface-name* ether-options],
[edit interfaces ge-*pim*/0/0 switch-options switch-port *port-number*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description Set the device's link connection characteristic.

Options *mode*—Link characteristics:

- **automatic**—Link mode is negotiated. This is the default for EX Series switches.
- **full-duplex**—Connection is full duplex.
- **half-duplex**—Connection is half duplex.

Default: Fast Ethernet interfaces can operate in either full-duplex or half-duplex mode. The router's or switch's management Ethernet interface, **fxp0** or **em0**, and the built-in Fast Ethernet interfaces on the FIC (M7i router) autonegotiate whether to operate in full-duplex or half-duplex mode. Unless otherwise noted here, all other interfaces operate only in full-duplex mode.

NOTE: On EX Series switches, if **no-auto-negotiation** is specified in **[edit interfaces *interface-name* ether-options]**, you can select only **full-duplex** or **half-duplex**. If **auto-negotiation** is specified, you can select any mode.

NOTE:

- Member links of an aggregated Ethernet bundle must not be explicitly configured with a link mode. You must remove any such link-mode configuration before committing the aggregated Ethernet configuration.
- Starting with Junos OS release 17.4R1 and later, the **link-mode** configuration is not supported for 10-Gigabit Ethernet interfaces.

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

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- *Configuring the Link Characteristics on Ethernet Interfaces*
- Understanding Management Ethernet Interfaces on page 10
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure) for EX Series Switches with ELS support*

 Copyright © 2018, Juniper Networks, Inc.

link-protection

Syntax

```
link–protection {
  disable;
  (revertive |non-revertive);
}
```

Hierarchy Level

[edit interfaces ae*x* aggregated-ether-options]
[edit interfaces ae*x* aggregated-ether-options *lacp*]

Release Information

Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.
Support for **disable**, **revertive**, and **non-revertive** statements added in Junos OS Release 9.3.

Description

On the router, for aggregated Ethernet interfaces only, configure link protection. In addition to enabling link protection, a primary and a secondary (backup) link must be configured to specify what links egress traffic should traverse. To configure primary and secondary links on the router, include the **primary** and **backup** statements at the **[edit interfaces ge-*fpc/pic/port* gigether-options 802.3ad ae*x*]** hierarchy level or the **[edit interfaces fe-*fpc/pic/port* fastether-options 802.3ad ae*x*]** hierarchy level.

On the switch, you can configure either Junos OS link protection for aggregated Ethernet interfaces or the LACP standards link protection for aggregated Ethernet interfaces.

For Junos OS link protection, specify **link-protection** at the following hierarchy levels:

- **[edit interfaces ge-*fpc/pic/port* ether-options 802.3ad ae*x*]**
- **[edit interfaces xe-*fpc/pic/port* ether-options 802.3ad ae*x*]** hierarchy level or at the **[edit interfaces xe-*fpc/pic/port* ether-options 802.3ad ae*x*]** hierarchy level.

  To disable link protection, use the **delete interface ae aggregate-ether-options link-protection** statement at the **[edit interfaces ae*x* aggregated-ether-options]** hierarchy level or the **[edit interfaces ae*x* aggregated-ether-options *lacp*]]** hierarchy level.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Configuring Aggregated Ethernet Link Protection*
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*

Copyright © 2018, Juniper Networks, Inc.

link-speed (Aggregated Ethernet)

Syntax link-speed *speed*;

Hierarchy Level (EX Series)
[edit interfaces aex aggregated-ether-options],
[edit interfaces interface-range *name* aggregated-ether-options],
[edit interfaces interface-range *name* aggregated-sonet-options]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
mixed option added in Junos OS Release 15.1F3 and 16.1R2 for PTX5000 routers and 15.1F6 and 16.1R2 for PTX3000 routers.

Description For aggregated Ethernet interfaces only, set the required link speed.

Options *speed*—For aggregated Ethernet links, you can specify ***speed*** in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Aggregated Ethernet links on the M120 router can have one of the following speeds:

- **100m**—Links are 100 Mbps.
- **10g**—Links are 10 Gbps.
- **1g**—Links are 1 Gbps.
- **oc192**—Links are OC192 or STM64c.

Aggregated Ethernet links on EX Series switches can be configured to operate at one of the following speeds:

- **10m**—Links are 10 Mbps.
- **100m**—Links are 100 Mbps.
- **1g**—Links are 1 Gbps.
- **10g**—Links are 10 Gbps.

Aggregated Ethernet links on T Series, MX Series, PTX Series routers, and QFX5100, QFX10002, QFX10008, and QFX10016 switches can be configured to operate at one of the following speeds:

- **100g**—Links are 100 Gbps.
- **100m**—Links are 100 Mbps.
- **10g**—Links are 10 Gbps.
- **1g**—Links are 1 Gbps.
- **40g**—Links are 40 Gbps.

 Copyright © 2018, Juniper Networks, Inc.

- **50g**—Links are 50 Gbps.
- **80g**—Links are 80 Gbps.
- **8g**—Links are 8 Gbps.
- **mixed**—Links are of various speeds.
- **oc192**—Links are OC192.

mixed—Enables bundling of different Ethernet rate links in the same Aggregated Ethernet interface.

Required Privilege Level

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

Related Documentation

- *Aggregated Ethernet Interfaces Overview*
- Configuring Aggregated Ethernet Link Speed on page 118
- *Configuring Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles*
- *Configuring Aggregated Ethernet Links (CLI Procedure)*
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*

Copyright © 2018, Juniper Networks, Inc.

link-speed (Aggregated SONET/SDH)

Syntax link-speed (*speed* | mixed);

Hierarchy Level [edit interfaces as*x* aggregated-sonet-options]

Release Information Statement introduced before Junos OS Release 7.4.
mixed option added in Release 8.0.

Description For aggregated SONET/SDH interfaces only, set the required link speed.

Options *speed*—Aggregated SONET/SDH links can have one of the following speed values.

- **oc3**—Links are OC3c or STM1c.
- **oc12**—Links are OC12c or STM4c.
- **oc48**—Links are OC48c or STM16c.
- **oc192**—Links are OC192c or STM64c.
- **oc768**—Links are OC768c or STM256c.

mixed—For aggregated SONET/SDH links on T Series routers, you can mix interface speeds in SONET/SDH aggregation bundles. Interface speeds from OC3 through OC768 are supported.

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

Related Documentation

- Configuring Aggregated Ethernet Link Speed on page 118
- *Configuring Aggregated SONET/SDH Interfaces*

 Copyright © 2018, Juniper Networks, Inc.

linktrace

Syntax

```
linktrace {
  age (30m | 10m | 1m | 30s | 10s);
  path-database-size path-database-size;
}
```

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure connectivity fault management linktrace parameters.

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

Related Documentation
- *Configuring Linktrace Protocol in CFM*

Copyright © 2018, Juniper Networks, Inc.

lmi (Frame Relay)

Syntax

```
lmi {
  lmi-type (ansi | itu | c-lmi);
  n391dte number;
  n392dce seconds;
  n392dte number;
  n393dce number;
  n393dte number;
  t391dte number;
  t392dce seconds;
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Set Frame Relay keepalive parameters.

Options

n391dte—DTE full status polling interval.

Range: 1 through 255

Default: 6

n392dce—DCE error threshold, in number of errors.

Range: 1 through 10

Default: 3

n392dte—DTE error threshold, in number of errors.

Range: 1 through 10

Default: 3

n393dce—DCE monitored event-count.

Range: 1 through 10

Default: 4

n393dte—DTE monitored event-count.

Range: 1 through 10

Default: 4

t391dte—DTE polling timer.

Range: 5 through 30 seconds

Default: 10 seconds

t392dce—DCE polling timer.

Range: 5 through 30 seconds

Default: 15 seconds

The remaining statements are explained separately. See CLI Explorer.

Copyright © 2018, Juniper Networks, Inc.

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

Related Documentation

- *Configuring Tunable Keepalives for Frame Relay LMI*
- lmi-type on page 741
- mlfr-uni-nni-bundle-options on page 779

lmi (Ethernet OAM)

Syntax

```
lmi {
  status-counter count;
  polling-verification-timer value;
  interface name {
    uni-id uni-name;
    status-counter number;
    polling-verification-timer value;
    evc-map-type (all-to-one-bundling | bundling | service-multiplexing);
    evc evc-name {
      default-evc;
      vlan-list vlan-id-list;
    }
  }
}
```

Hierarchy Level [edit protocols oam ethernet]

Release Information Statement introduced in Junos OS Release 9.5.

Description On routers with **ge**, **xe**, or **ae** interfaces, configure an OAM Ethernet Local Management Interface (E-LMI).

NOTE: On MX Series routers, E-LMI is supported on Gigabit Ethernet (ge), 10-Gigabit Ethernet (xe), and Aggregated Ethernet (ae) interfaces configured on MX Series routers with DPC only.

Options

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

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

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

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

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

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

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

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

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

Copyright © 2018, Juniper Networks, Inc.

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

Related Documentation
- *Configuring Ethernet Local Management Interface*
- evcs on page 599

lmi-type

Syntax lmi-type (ansi | itu | c-lmi);

Hierarchy Level
[edit interfaces *interface-name* lmi],
[edit interfaces *interface-name* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Set Frame Relay Local Management Interface (LMI) type.

NOTE: Consortium LMI is supported on all MPCs and I-chip based FPCs.

Options
ansi—Use ANSI T1.617 Annex D LMIs.

itu—Use ITU Q933 Annex A LMIs.

c-lmi—Use Consortium LMI.
Default: ansi

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

Related Documentation
- *Configuring Frame Relay Keepalives*
- mlfr-uni-nni-bundle-options on page 779
- lmi (Frame Relay) on page 738
- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2018, Juniper Networks, Inc.

load-interval

Syntax load-interval *seconds*;

Hierarchy Level [edit interfaces dl*n* unit *logical-unit-number* dialer-options],
[edit logical-systems *logical-system-name* interfaces dl*n* unit *logical-unit-number* dialer-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description On J Series Services Routers with ISDN logical interfaces, specify the interval used to calculate the average load on the network. By default, the average interface load is calculated every 60 seconds.

Options ***seconds***—Number of seconds at which the average load calculation is triggered.

Range: 20 through 180, in 10-second intervals

Default: 60 seconds

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

Related Documentation

- *Junos OS Interfaces and Routing Configuration Guide*

 Copyright © 2018, Juniper Networks, Inc.

load-threshold

Syntax load-threshold *percent*;

Hierarchy Level [edit interfaces dl*n* unit *logical-unit-number* dialer-options],
[edit logical-systems *logical-system-name* interfaces dl*n* unit *logical-unit-number* dialer-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description On J Series Services Routers with ISDN logical interfaces, specify the bandwidth threshold percentage used for adding interfaces. Another link is added to the multilink bundle when the load reaches the threshold value you set. Specify a percentage between 0 and 100.

Options *percent*—Bandwidth threshold percentage used for adding interfaces. When set to 0, all available channels are dialed.

Range: 0 through 100 seconds

Default: 100 seconds

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

Related Documentation

- *Junos OS Interfaces and Routing Configuration Guide*

Copyright © 2018, Juniper Networks, Inc.

local-name

Syntax local-name *name*;

Hierarchy Level [edit interfaces *interface-name* ppp-options chap],
[edit interfaces *interface-name* ppp-options pap],
[edit interfaces *interface-name* unit *logical-unit-number* ppp-options chap],
[edit interfaces *interface-name* unit *logical-unit-number* ppp-options pap],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* ppp-options chap],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* ppp-options pap]
[edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" ppp-options],
[edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit" ppp-options]

Release Information Statement introduced before Junos OS Release 7.4.
Support for PAP added in Junos OS Release 8.3.
Support at the **[edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit" ppp-options]** hierarchy level introduced in Junos OS Release 14.2.

Description Specify the name of the interface used for CHAP or PAP authentication. Dynamic interfaces are supported only for CHAP authentication.

For ATM2 IQ interfaces only, you can configure a CHAP local name on the logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:

- **atm-ppp-llc**—PPP over AAL5 LLC encapsulation.
- **atm-ppp-vc-mux**—PPP over AAL5 multiplex encapsulation.

Options ***name***—Name of the interface used as an identifier in CHAP challenge and response packets or PAP request and response packets.

Default: When you do not include the **local-name** statement in the configuration, the interface sends the router's system hostname in CHAP challenge and response packets or PAP request and response packets.

Range: For CHAP authentication, a string of 1 through 32 characters. For PAP authentication, a string of 1 through 8 characters.

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

Related Documentation

- Configuring the PPP Challenge Handshake Authentication Protocol on page 136
- Configuring the PPP Password Authentication Protocol On a Physical Interface on page 139

 Copyright © 2018, Juniper Networks, Inc.

- *Junos OS Administration Library*

local-password

Syntax local-password *password*;

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

Release Information Statement introduced in Junos OS Release 8.3.

Description Configure the host password for sending PAP requests.

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

Related Documentation

- *Configuring the Local Password*
- Configuring the PPP Password Authentication Protocol On a Physical Interface on page 139

lockout

Syntax lockout;

Hierarchy Level [edit interfaces *interface-name* sonet-options aps]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure a lockout of protection, forcing the use of the working circuit and locking out the protect circuit regardless of anything else.

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

Related Documentation

- *Configuring Switching Between the Working and Protect Circuits*

Copyright © 2018, Juniper Networks, Inc.

log-prefix (Interfaces)

Syntax log-prefix *prefix-value*;

Hierarchy Level [edit interfaces *interface-name* services-options syslog host *hostname*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Set the system logging prefix value.

Options ***prefix-value***—System logging prefix value.

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*
- *Configuring System Logging for Services Interfaces*

logical-interface-fpc-redundancy (Aggregated Ethernet Subscriber Interfaces)

Syntax logical-interface-fpc-redundancy;

Hierarchy Level [edit interfaces ae*number* aggregated-ether-options]

Release Information Statement introduced in Junos OS Release 11.2.
Statement introduced in Junos OS Release 13.2R2 for EX Series switches.

Description Provide module redundancy for demux subscribers on aggregated Ethernet bundles configured with targeted distribution. Backup links for a subscriber are chosen on a different EQ DPC or MPC from the primary link, based on the link with the fewest number of subscribers among the links on different modules. If all links are on a single module when this is configured, backup links are not provisioned.

By default, link redundancy is provided for the aggregated Ethernet bundle.

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

Related Documentation

- *Configuring Link and Module Redundancy for Demux Subscribers in an Aggregated Ethernet Interface*
- *Configuring Module Redundancy for a Virtual Chassis*

 Copyright © 2018, Juniper Networks, Inc.

logical-interface-policer

Syntax logical-interface-policer;

Hierarchy Level [edit dynamic-profiles *profile-name* firewall policer *policer-name*],
[edit dynamic-profiles *profile-name* firewall three-color-policer *name*],
[edit firewall atm-policer*atm-policer-name*],
[edit firewall policer *policer-name*],
[edit firewall policer *policer-template-name*],
[edit firewall three-color-policer *policer-name*],
[edit logical-systems *logical-system-name* firewall policer *policer-name*],
[edit logical-systems *logical-system-name* firewall three-color-policer *name*]

Release Information Statement introduced before Junos OS Release 7.4.
Support at the **[edit firewall three-color-policer *policer-name*]** hierarchy level introduced in Junos OS Release 8.2.
Logical systems support introduced in Junos OS Release 9.3.
Support at the **[edit dynamic-profiles ... policer *policer-name*]** and **[edit dynamic-profiles ... three-color-policer *name*]**hierarchy levels introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure a logical interface policer.

NOTE: Starting in Junos OS Release 12.2R2, on T Series Core Routers only, you can configure an MPLS LSP policer for a specific LSP to be shared across different protocol family types. You must include the **logical-interface-policer** statement to do so.

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

Related Documentation

- *Two-Color and Three-Color Logical Interface Policers*
- *Traffic Policer Types*
- *Configuring and Applying Tricolor Marking Policers*
- *action*
- *Configuring Gigabit Ethernet Two-Color and Tricolor Policers*
- action on page 418

Copyright © 2018, Juniper Networks, Inc.

logical-systems

Syntax

```
logical-systems {
  logical-system-name {
    ...logical-system-configuration...
  }
}
```

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 7.4.
Statement name changed from **logical-routers** in Junos OS Release 9.3.

Description Configure a logical system.

Options ***logical-system-name***—Name of the logical system.

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

Related Documentation

- *Logical Systems Feature Guide for Routers and Switches*

 Copyright © 2018, Juniper Networks, Inc.

long-buildout

Syntax (long-buildout | no-long-buildout);

Hierarchy Level [edit interfaces *interface-name* t3-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the T3 line buildout. A T3 interface has two settings for the T3 line buildout: a short setting, which is less than 255 feet (68 meters), and a long setting, which is greater than 255 feet and shorter than 450 feet (137 meters).

This statement applies to copper-cable-based T3 interfaces only. You cannot configure a line buildout for a DS3 channel on a channelized OC12 interface, which runs over fiber-optic cable.

Default A T3 interface uses the short line buildout setting (**no-long-buildout**) for wires shorter than 255 feet (68 meters).

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

Related Documentation

- *Configuring the T3 Line Buildout*

loop-timing

Syntax

(loop-timing | no-loop-timing);

Hierarchy Level

[edit interfaces ct3-*fpc/pic/port* t3-options],
[edit interfaces e1-*fpc/pic/port*:0 sonet-options],
[edit interfaces stm1-*fpc/pic/port* sonet-options]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

For channelized IQ interfaces and non-IQ channelized STM1 interfaces only, configure the SONET/SDH or DS3-level clocking source.

NOTE: On M Series, MX Series, and T Series routers, under E1 channels, loop timing can be configured only at channel 0. When you configure on channel 0, it is applicable on all channels as internal by default.

Options

loop-timing—Configure loop timing (external) clocking.

no-loop-timing—Configure line timing (internal) clocking.

Default: no-loop-timing

Required Privilege Level

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

Related Documentation

- *Configuring Channelized IQ and IQE SONET/SDH Loop Timing*
- *Configuring the Channelized T3 Loop Timing*
- clocking on page 486

Copyright © 2018, Juniper Networks, Inc.

loopback (ADSL, DS0, E1/E3, SONET/SDH, SHDSL, and T1/T3)

Syntax loopback (local | payload | remote);

Hierarchy Level [edit interfaces ce1-*fpc/pic/port*],
[edit interfaces ct1-*fpc/pic/port*],
[edit interfaces t1-*fpc/pic/port*],
[edit interfaces *interface-name* ds0-options],
[edit interfaces *interface-name* dsl-options],
[edit interfaces *interface-name* e1-options],
[edit interfaces *interface-name* e3-options],
[edit interfaces *interface-name* shdsl-options},
[edit interfaces *interface-name* sonet-options],
[edit interfaces *interface-name* t1-options],
[edit interfaces *interface-name* t3-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers.

Description Configure a loopback connection. To turn off the loopback capability, remove the **loopback** statement from the configuration.

NOTE: When configuring CE1 or CT1 interfaces on 10-port Channelized E1/T1 IQE PICs, the **loopback** statement must be included with the **local** or **remote** option at the **[edit interfaces ce1-*fpc/pic/port*]** or **[edit interfaces ct1-*fpc/pic/port*]** hierarchy level as appropriate.

When configuring T1 interfaces on 10-port Channelized E1/T1 IQE PICs, the **loopback** statement must be included with the **payload** option at the **[edit interfaces t1-*fpc/pic/port*]** hierarchy level.

NOTE: When configuring CE1 or CT1 interfaces on the 16-port Channelized E1/T1 MIC (MIC-3D-16CHE1-T1-CE), you must include the **loopback** statement at the **[edit interfaces ce1-*fpc/pic/port*]** hierarchy level, or **[edit interfaces ct1-*fpc/pic/port*]**

To configure loopback on channelized IQ and IQE PICs, SONET/SDH level, use the **sonet-options loopback** statement **local** and **remote** options at the controller interface (coc48, cstm16, coc12, cstm4, coc3, cstm1). It is ignored for path-level interfaces **so-*fpc/pic/port*** or **so-*fpc/pic/port*:*channel***.

Options **local**—Loop packets, including both data and timing information, back on the local router's PIC. *N*xDS0 IQ interfaces do not support local loopback.

Copyright © 2018, Juniper Networks, Inc.

payload—For channelized T3, T1, and *N*xDS0 IQ interfaces only, loop back data only (without clocking information) on the remote router's PIC. With payload loopback, overhead is recalculated. Neither ATM-over-asymmetrical digital subscriber line (ADSL) interfaces nor ATM-over-SHDSL interfaces support payload loopback.

remote—Loop packets, including both data and timing information, back on the remote router's interface card. *N*xDS0 IQ interfaces do not support remote loopback.

Required Privilege Level

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

Related Documentation

- *Configuring E3 and T3 Parameters on ATM Interfaces*
- *Configuring E1 Loopback Capability*
- *Configuring E3 Loopback Capability*
- *Configuring SONET/SDH Loopback Capability to Identify a Problem as Internal or External*
- *Configuring SHDSL Operating Mode on an ATM Physical Interface*
- *Configuring T1 Loopback Capability*
- *Configuring T3 Loopback Capability*
- feac-loop-respond on page 614

 Copyright © 2018, Juniper Networks, Inc.

loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet)

Syntax (loopback | no-loopback);

Hierarchy Level [edit interfaces *interface-name* aggregated-ether-options],
[edit interfaces *interface-name* ether-options],
[edit interfaces *interface-name* fastether-options],
[edit interfaces *interface-name* gigether-options],
[edit interfaces interface-range *name* ether-options]

For QFX Series and EX Series:

[edit interfaces *interface-name* aggregated-ether-options],
[edit interfaces *interface-name* ether-options],

For SRX Series Devices and vSRX:

[edit interfaces *interface-name* redundant-ether-options]

Release Information Statement introduced before Junos OS Release 7.4 for MX Series.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Statement modified in Junos OS Release 9.2 for the SRX Series.

Description For aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces, enable or disable loopback mode.

NOTE:

- By default, local aggregated Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces connect to a remote system.
- IPv6 Neighbor Discovery Protocol (NDP) addresses are not supported on Gigabit Ethernet interfaces when loopback mode is enabled on the interface. That is, if the **loopback** statement is configured at the **[edit interfaces ge-*fpc/pic/port* gigether-options]** hierarchy level, an NDP address cannot be configured at the **[edit interfaces ge-*fpc/pic/port* unit *logical-unit-number* family inet6 address]** hierarchy level.

Default By default, loopback is disabled.

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

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- *Configuring Ethernet Loopback Capability*
- *Understanding Interfaces*

loopback (Serial)

Syntax loopback *mode*;

Hierarchy Level [edit interfaces *interface-name* serial-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure a loopback connection.

Default If you do not include this statement, there is no loopback connection.

Options *mode*—You can specify the one of the following loopback modes:

- **dce-local**—For EIA-530 interfaces only, loop packets back on the local DCE.
- **dce-remote**—For EIA-530 interfaces only, loop packets back on the remote DCE.
- **local**—Loop packets back on the local router's PIC.
- **remote**—Loop packets back on the line interface unit (LIU).

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

Related Documentation

- Configuring Serial Loopback Capability on page 348

 Copyright © 2018, Juniper Networks, Inc.

loopback-clear-timer

Syntax loopback-clear-timer *seconds*;

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

Release Information Statement introduced in Junos OS Release 8.5.

Description For interfaces with PPP, PPP TCC, PPP over Ethernet, PPP over ATM, and PPP over Frame Relay encapsulations, configure a loop detection clear timer for the Link Control Protocol (LCP) component of a PPP session.

Options ***seconds***—The time in seconds to wait before the loop detection flag is cleared if it is not cleared by the protocol.

Range: 1 through 60 seconds

Default: 9 seconds

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

Related Documentation

- Configuring the PPP Clear Loop Detected Timer on page 194

Copyright © 2018, Juniper Networks, Inc.

loss-priority

Syntax loss-priority (high | low);

Hierarchy Level [edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile output-priority-map classifier premium forwarding-class *class-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the packet loss priority value.

Options **high**—Packet has high loss priority.

low—Packet has low loss priority.

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

Related Documentation

- *Specifying an Output Priority Map*

 Copyright © 2018, Juniper Networks, Inc.

loss-threshold

Syntax loss-threshold *number*;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* continuity-check]

Release Information Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

Description Specify the number of continuity check messages lost before marking the remote MEP as down. The value can be from 3 to 256 protocol data units (PDUs). The default value is 3 PDUs.

Options *number*—The number of continuity check messages that can be lost before the remote MEP is considered down.

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

Related Documentation

- *Continuity Check Protocol Parameters Overview*
- *Configuring Continuity Check Protocol Parameters for Fault Detection*

Copyright © 2018, Juniper Networks, Inc.

low-plp-max-threshold

Syntax low-plp-max-threshold *percent*;

Hierarchy Level [edit interfaces at-*fpc/pic/port* atm-options linear-red-profiles *profile-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For ATM2 IQ interfaces only, define the drop profile fill-level for the low PLP CoS VC. When the fill level exceeds the defined percentage, all packets are dropped.

Options ***percent***—Fill-level percentage when linear RED is applied to cells with PLP.

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

Related Documentation

- *ATM2 IQ VC Tunnel CoS Components Overview*
- high-plp-max-threshold on page 640
- low-plp-threshold on page 759
- *Configuring Linear RED Profiles on ATM Interfaces*
- *high-plp-max-threshold*
- queue-depth on page 914

 Copyright © 2018, Juniper Networks, Inc.

low-plp-threshold

Syntax low-plp-threshold *percent*;

Hierarchy Level [edit interfaces at-*fpc/pic/port* atm-options linear-red-profiles *profile-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For ATM2 IQ interfaces only, define the CoS VC drop profile fill-level percentage when linear RED is applied to cells with low PLP. When the fill level exceeds the defined percentage, packets with low PLP are randomly dropped by RED. This statement is mandatory.

Options ***percent***—Fill-level percentage when linear RED is applied to cells with low PLP.

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

Related Documentation

- *ATM2 IQ VC Tunnel CoS Components Overview*
- high-plp-max-threshold on page 640
- high-plp-threshold on page 636
- *Configuring Linear RED Profiles on ATM Interfaces*
- *high-plp-max-threshold*
- *high-plp-threshold*
- low-plp-max-threshold on page 758
- queue-depth on page 914

Copyright © 2018, Juniper Networks, Inc.

lowest-priority-defect

Syntax lowest-priority-defect (all-defects | err-xcon | mac-rem-err-xcon | no-defect | rem-err-xcon | xcon)

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id*]

Release Information Statement introduced in Junos OS Release 10.0.

Description Specify the lowest priority defect that is allowed to generate a Fault Alarm whenever CFM detects a defect. This configuration is done at the MEP level.

Options Specify one of the following lowest priority defect options:

all-defects—Allows all defects.

err-xcon—Allows only erroneous CCM and cross-connect CCM defects.

mac-rem-err-xcon—Allows only MAC, not receiving CCM, erroneous CCM, and cross-connect defects.

no-defect—Allows no defects.

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

xcon—Allows only cross-connect CCM defects.

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

Related Documentation
- *Configuring the Maintenance End Point Lowest Priority Defect*

 Copyright © 2018, Juniper Networks, Inc.

lsq-failure-options

Syntax

```
lsq-failure-options {
  no-termination-request;
  [trigger-link-failure interface-name];
}
```

Hierarchy Level [edit interfaces lsq-*fpc/pic/port*]

Release Information Statement introduced in Junos OS Release 7.4.

Description For AS PIC or MultiServices PIC link services IQ (**lsq**) interfaces only, define the failure recovery option settings.

Options The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

mac

Syntax mac *mac-address*;

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Set the MAC address of the interface.

Use this statement at the **[edit interfaces ... ps0]** hierarchy level to configure the MAC address for a pseudowire logical device that is used for subscriber interfaces over point-to-point MPLS pseudowires.

Options ***mac-address***—MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: ***nnnn.nnnn.nnnn*** or ***nn:nn:nn:nn:nn:nn***. For example, **0000.5e00.5355** or **00:00:5e:00:53:55**.

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

Related Documentation

- *Configuring the MAC Address on the Management Ethernet Interface*
- *Configuring a Pseudowire Subscriber Logical Interface Device*

Copyright © 2018, Juniper Networks, Inc.

mac-address (Accept Source Mac)

Syntax
mac-address *mac-address* policer;

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

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description
For Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), for Gigabit Ethernet DPCs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP, specify a remote MAC address on which to count incoming and outgoing packets.

Options
mac-address—MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: ***nnnn.nnnn.nnnn*** or ***nn:nn:nn:nn:nn:nn***. For example, **0011.2233.4455** or **00:11:22:33:44:55**.

policer—MAC policer. For more information, see policer (MAC).

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

Related Documentation
- *Configuring MAC Address Filtering*

mac-address (VLAN and Stacked VLAN Interfaces)

Syntax
mac-address;

Hierarchy Level
[edit interfaces *interface-name* auto-configure vlan-ranges authentication username-include],
[edit interfaces *interface-name* auto-configure stacked-vlan-ranges authentication username-include],

Release Information
Statement introduced in Junos OS Release 10.0.

Description
Specify that the client hardware address (chaddr) from the incoming DHCP discover packet be concatenated with the username during the subscriber authentication process.

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

Related Documentation
- *Configuring VLAN Interface Username Information for AAA Authentication*

Copyright © 2018, Juniper Networks, Inc.

mac-learn-enable

Syntax

mac-learn-enable;

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile]
[edit interfaces ae*x* aggregated-ether-options ethernet-switch-profile]

Release Information

Statement introduced before Junos OS Release 7.4.
Support for statement under the **[edit interfaces ae*x* aggregated-ether-options ethernet-switch-profile]** hierarchy introduced in Junos OS Release 15.1.

Description

For Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), for Gigabit Ethernet DPCs on MX Series routers, for 100-Gigabit Ethernet Type 5 PIC with CFP, and for MPC3E, MPC4E, MPC5E, MPC5EQ, and MPC6E MPCs, configure dynamic learning of the source and destination MAC addresses. By default, the interface is not allowed to dynamically learn source and destination MAC addresses.

To disable dynamic learning of the source and destination MAC addresses after it has been configured, you must delete **mac-learn-enable** from the configuration.

MPCs support MAC address accounting for an individual interface or an aggregated Ethernet interface member link only after the interface has received traffic from the MAC source. If traffic is only exiting an interface, the MAC address is not learned and MAC address accounting does not occur.

Required Privilege Level

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

Related Documentation

- *Configuring MAC Address Filtering*
- *Configuring MAC Address Accounting*

 Copyright © 2018, Juniper Networks, Inc.

mac-validate

Syntax mac-validate (loose | strict);

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

Release Information Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Enable IP and MAC address validation for static Ethernet and IP demux interfaces.

Options **loose**—Forwards incoming packets when both the IP source address and the MAC source address match one of the trusted address tuples. Drops packets when the IP source address matches one of the trusted tuples, but the MAC address does not match the MAC address of the tuple. Continues to forward incoming packets when the source address of the incoming packet does not match any of the trusted IP addresses.

strict—Forwards incoming packets when both the IP source address and the MAC source address match one of the trusted address tuples. Drops packets when the MAC address does not match the tuple's MAC source address, or when IP source address of the incoming packet does not match any of the trusted IP addresses.

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

Related Documentation

- *MAC Address Validation on Static Ethernet Interfaces Overview*
- Configuring an IP Demultiplexing Interface on page 316
- Configuring a VLAN Demultiplexing Interface on page 320

maintenance-association

Syntax

```
maintenance-association ma-name {
  short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
  protect-maintenance-association protect-ma-name;
  remote-maintenance-association remote-ma-name;
  continuity-check {
    hold-interval minutes;
    interval (10m | 10s | 1m | 1s| 100ms);
    loss-threshold number;
  }
  mep mep-id {
    auto-discovery;
    direction (up | down);
    interface interface-name (protect | working);
    lowest-priority-defect (all-defects | err-xcon | mac-rem-err-xcon | no-defect |
      rem-err-xcon | xcon );
    priority number;
    remote-mep mep-id {
      action-profile profile-name;
      sla-iterator-profile profile-name {
        data-tlv-size size;
        iteration-count count-value;
        priority priority-value;
      }
    }
  }
}
```

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name*]

Release Information Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

Description Configure the name of the maintenance association in IEEE-compliant format.

Options **ma-name**—The name of the maintenance association within the maintenance domain.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *Creating a Maintenance Association*
- *Configuring a MEP to Generate and Respond to CFM Protocol Messages*

 Copyright © 2018, Juniper Networks, Inc.

maintenance-domain

Syntax

```
maintenance-domain domain-name {
  bridge-domain name <vlan-id [ vlan-ids ]>;
  instance vpls-instance-name;
  level number;
  maintenance-association ma-name {
    protect-maintenance-association  protect-ma-name;
    remote-maintenance-association  remote-ma-name;
    short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
    continuity-check {
      hold-interval minutes;
      interval (10m | 10s | 1m | 1s| 100ms);
      loss-threshold number
    }
    mep mep-id {
      auto-discovery;
      direction (up | down);
      interface interface-name (protect | working);
      lowest-priority-defect (all-defects | err-xcon | mac-rem-err-xcon | no-defect |
        rem-err-xcon | xcon );
      priority number;
      remote-mep mep-id {
        action-profile profile-name;
        sla-iterator-profile profile-name {
          data-tlv-size size;
          iteration-count count-value;
          priority priority-value;
        }
      }
    }
    mip-half-function(none | default | explicit);
    name-format (character-string | none | dns | mac+2oct);
  }
  virtual-switch name {
    bridge-domain name <vlan-id [ vlan-ids ]>;
  }
}
```

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management]

Release Information Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.
Support for multiple down MEP introduced in Junos OS Release 15.1R1 for MX Series Routers.

Description Configure the name of the maintenance domain in IEEE-compliant format.

NOTE: For MX Series Routers, you can configure multiple down MEPs for a single instance of maintenance domain identifier and maintenance

Copyright © 2018, Juniper Networks, Inc.

association name to monitor services provided on Virtual Private LAN Service (VPLS), bridge, circuit cross-connect (CCC), and IPv4 domains.

Options *domain-name*—Name of the maintenance domain.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *Creating a Maintenance Domain*
- *Configuring a MEP to Generate and Respond to CFM Protocol Messages*

master-only

Syntax master-only;

Hierarchy Level [edit groups re*x* interfaces (fxp0 | em0) unit *logical-unit-number* family *family* address],
[edit groups re*x* logical-systems *logical-system-name* interfaces fxp0 unit *logical-unit-number* family *family* address],
[edit interfaces (fxp0 | em0) unit *logical-unit-number* family *family* address],
[edit logical-systems *logical-system-name* interfaces fxp0 unit *logical-unit-number* family *family* address]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure the IP address to be used when the Routing Engine is the current master.

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

Related Documentation

- *Configuring a Consistent Management IP Address*
- *CLI User Guide*

 Copyright © 2018, Juniper Networks, Inc.

maximum-contexts

Syntax maximum-contexts *number* <force>;

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

Release Information Statement introduced in Junos OS Release 7.5.

Description Specify the maximum number of RTP contexts to accept during negotiation.

Options ***number***—Maximum number of contexts.

force—(Optional) Requires the PIC to use the value specified for maximum RTP contexts, regardless of the negotiated value. This option allows the software to interoperate with Junos OS Releases that base the RTP context value on link speed.

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

maximum-requests

Syntax maximum-requests *times*;

Hierarchy Level [edit protocols dot1x authenticator interface *interface-id*]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the maximum number of retransmission times of an EAPOL Request packet to the client before it times out the authentication session.

Options *times*—Specify the maximum number of retransmission times.

Range: 1 through 10 times

Default: 2 times

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

Related Documentation

- *IEEE 802.1x Port-Based Network Access Control Overview*
- authenticator on page 448
- dot1x on page 543
- interface (IEEE 802.1x) on page 682

 Copyright © 2018, Juniper Networks, Inc.

maximum-vcs

Syntax maximum-vcs *maximum-vcs*;

Hierarchy Level [edit interfaces at-*fpc/pic/port* atm-options vpi *vpi-identifier*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For ATM1 interfaces, configure the maximum number of virtual circuits (VCs) allowed on a virtual path (VP). When configuring ATM1 interfaces on the router, you must include this statement.

For a configured virtual path identifier (VPI), valid virtual channel identifier (VCI) numbers are from 0 through (***maximum-vcs*** value – 1). VCI numbers 0 through 31 are reserved by the ATM Forum. It is recommended that you use a VCI number higher than 31 when connecting to an ATM switch.

Options *maximum-vcs*—Maximum number of VCs on the VP.

Range: 1 through 4090

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

Related Documentation

- *Configuring the Maximum Number of ATM1 VCs on a VP*
- multipoint-destination on page 794
- promiscuous-mode on page 905
- vci on page 1073

Copyright © 2018, Juniper Networks, Inc.

mc-ae

Syntax

```
mc-ae {
  chassis-id chassis-id;
    events {
      iccp-peer-down;
        force-icl-down;
      prefer-status-control-active;
      }
      init-delay-time seconds;
      mc-ae-id mc-ae-id;
      mode (active-active | active-standby);
      redundancy-group group-id;
      revert-time revert-time;
      status-control (active | standby);
      switchover-mode (non-revertive |revertive);
    }
```

Hierarchy Level

[edit interfaces ae*X* aggregated-ether-options],
[edit logical-systems *logical-system-name interfaces* ae*X* aggregated-ether-options]

Release Information

Statement introduced in Junos OS Release 9.6 for MX Series routers.
events statement introduced in Junos OS Release 11.4R4 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series. Only the **chassis-id**, **mc-ae-id**, **mode active-active**, and **status-control (active | standby)** options are supported on QFX Series devices.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
prefer-status-control-active statement introduced in Junos OS Release 13.2R1 for EX Series switches.
init-delay-time *seconds* statement introduced in Junos OS Release 13.2R3 for EX Series switches.
switchover-mode and **revert-time** statements introduced in Junos OS Release 13.3.
Support for logical systems introduced in Junos OS Release 14.1.

Description

Enable multichassis link aggregation groups (MC-LAG), which enables one device to form a logical LAG interface with two or more other devices.

Options

chassis-id—Specify the chassis ID for Link Aggregation Control Protocol (LACP) to calculate the port number of MC-LAG physical member links.

Values: 0 or 1

events—Specify an action if a specific MC-LAG event occurs.

iccp-peer-down—Specify an action if the ICCP peer of this node goes down.

force-icl-down—If the node's ICCP peer goes down, bring down the interchassis-link logical interface.

 Copyright © 2018, Juniper Networks, Inc.

prefer-status-control-active—Specify that the node configured as **status-control active** become the active node if the peer of this node goes down.

When ICCP goes down, you can use this keyword to make a mc-lag PE to become the active PE. For example, if you want mc-lag PE1 to be Active on ICCP down, then configure this keyword in PE1. It is not recommended to configure this keyword in both the mc-lag PEs.

NOTE: The **prefer-status-control-active** statement can be configured with the **status-control standby** configuration to prevent the LACP MC-LAG system ID from reverting to the default LACP system ID on ICCP failure. Use this configuration only if you can ensure that ICCP will not go down unless the router or switch is down. You must also configure the **hold-time down** value (at the **[edit interfaces *interface-name*]** hierarchy level) for the interchassis link with the **status-control standby** configuration to be higher than the ICCP BFD timeout. This configuration prevents data traffic loss by ensuring that when the router or switch with the **status-control active** configuration goes down, the router or switch with the **status-control standby** configuration does not go into standby mode.

To make the **prefer-status-control-active** configuration work with the **status-control standby** configuration when an interchassis-link logical interface is configured on aggregate Ethernet interface, you must either configure the **lacp periodic *interval*** statement at the **[edit interface *interface-name* aggregated-ether-options]** hierarchy level as **slow** or configure the **detection-time threshold** statement at the **[edit protocols iccp peer liveness-detection]** hierarchy level as less than 3 seconds.

init-delay-time *seconds*—To minimize traffic loss, specify the number of seconds in which to delay bringing the multichassis aggregated Ethernet interface back to the up state when you reboot an MC-LAG peer.

mc-ae-id *mc-ae-id*—Specify the identification number of the MC-LAG device. The two MC-LAG network devices that manage a given MC-LAG must have the same identification number.

Range: 1 through 65,535

mode (active-active | active-standby)—Specify whether the MC-LAG is in active-active or active-standby mode.

NOTE: You can configure IPv4 (inet) and IPv6 (inet6) addresses on **mc-ae** interfaces when the **active-standby** mode is configured.

Copyright © 2018, Juniper Networks, Inc.

redundancy-group ***group-id***—Specify the redundancy group identification number. The Inter-Chassis Control Protocol (ICCP) uses the redundancy group ID to associate multiple chassis that perform similar redundancy functions.

Range: 1 through 4,294,967,294

revert-time—Wait interval (in minutes) before the switchover to the preferred node is performed when the **switchover-mode** is configured as revertive.

Range: 1 through 10

status-control (active | standby)—Specify whether the chassis becomes active or remains in standby mode when an interchassis link failure occurs.

switchover-mode (non-revertive | revertive)—Specify whether Junos OS should trigger a link switchover to the preferred node when the active node is available.

NOTE: For **revertive** mode to automatically switch over to the preferred node, the **status-control** statement should be configured as **active**.

init-delay-time ***seconds***—To minimize traffic loss, specify the number of seconds by which to delay bringing the multichassis aggregated Ethernet (mc-ae) interface back to the up state when you reboot an MC-LAG peer.

Required Privilege Level

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

Related Documentation

- *Active-Active Bridging and VRRP over IRB Functionality Overview*
- *Configuring Multichassis Link Aggregation on MX Series Routers*
- *Configuring Multichassis Link Aggregation on EX Series Switches*
- *Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation*
- *Example: Configuring Multichassis Link Aggregation in Active-Active Mode*
- *Configuring Manual and Automatic Link Switchover for MC-LAG Interfaces*

 Copyright © 2018, Juniper Networks, Inc.

member-interface-speed

Syntax member-interface-speed *speed*;

Hierarchy Level [edit interfaces container-options member-interface-type]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify container-interface member-interface speed options.

Options ***speed***—Set interface speed to OC3, OC12, OC48, OC192, OC768, or mixed.

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

Related Documentation

- *Displaying APS Using a Container Interface with ATM Encapsulation*
- *Configuring Container Interfaces for APS on SONET Links*
- container-options on page 499

member-interface-type

Syntax

```
member-interface-type sonet {
  member-interface-speed [ speed ];
}
```

Hierarchy Level [edit interfaces container-options]

Release Information Statement introduced in Junos OS Release 9.2.

Description Specify container-interface member-interface type as sonet and speed options.

Options **sonet**—Protocol type of the container interface, specify sonet.

speed—Set interface speed to OC3, OC12, OC48, OC192, OC768, or mixed.

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

Related Documentation

- *Displaying APS Using a Container Interface with ATM Encapsulation*
- *Configuring Container Interfaces for APS on SONET Links*
- container-options on page 499

Copyright © 2018, Juniper Networks, Inc.

mep

Syntax

```
mep mep-id {
  action-profile action-profile-name
  auto-discovery;
  direction (up | down);
  interface interface-name (protect | working);
  priority number;
  remote-mep mep-id {
    action-profile profile-name;
    sla-iterator-profile profile-name {
      data-tlv-size size;
      iteration-count count-value;
      priority priority-value;
    }
  }
}
```

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name*]

Release Information Statement introduced in Junos OS Release 8.4.

Description The numeric identifier of the maintenance association end point (MEP) within the maintenance association.

Options **mep** ***mep-id***—Specify the numeric identifier of the MEP.

Range: 1 through 8191

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *Ethernet Alarm Indication Signal (ETH-AIS) Function Overview*
- *Configuring ETH-AIS on a CFM MEP*
- *Configuring a MEP to Generate and Respond to CFM Protocol Messages*

 Copyright © 2018, Juniper Networks, Inc.

minimum-links

Syntax (SRX, MX, T, M, EX, QFX Series, EX4600, Qfabric System)

minimum-links *number*;

Hierarchy Level (EX Series)

[edit interfaces ae*x* aggregated-ether-options],
[edit interfaces ae*x* aggregated-sonet-options],
[edit interfaces *interface-name* mlfr-uni-nni-bundle-options],
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit interfaces interface-range *range* aggregated-ether-options],
[edit interfaces interface-range *range* aggregated-sonet-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Hierarchy Level (QFX Series)

[edit interfaces ae*x* aggregated-ether-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description

For aggregated Ethernet, SONET/SDH, multilink, link services, and voice services interfaces only, set the minimum number of links that must be up for the bundle to be labeled up.

Options

number—Number of links.

Range: On M120, M320, MX Series, T Series, and TX Matrix routers with Ethernet interfaces, the valid range for minimum-links number is 1 through 64. When the maximum value (16) is specified, all configured links of a bundle must be up for the bundle to be labeled up.On all other routers and on EX Series switches, other than EX8200 switches, the range of valid values for minimum-links number is 1 through 8. When the maximum value (8) is specified, all configured links of a bundle must be up for the bundle to be labeled 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.On EX4600, QFX Series and Q Fabric Systems , 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.

Default: 1

Required Privilege Level

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

Related Documentation

- *Configuring Aggregated Ethernet Minimum Links*
- *Configuring Aggregated SONET/SDH Interfaces*
- *Configuring Aggregated Ethernet Links (CLI Procedure)*

Copyright © 2018, Juniper Networks, Inc.

- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*
- *Junos OS Services Interfaces Library for Routing Devices*
- *Configuring Link Aggregation*

mip-half-function

Syntax mip-half-function (none | default | explicit);

Hierarchy Level [edit protocols oam ethernet connectivity-fault-managementmaintenance-domain *md-name*],
[edit protocols oam ethernet connectivity-fault-managementmaintenance-association *ma-name*]

Release Information Statement introduced in Junos OS Release 9.6.

Description Specify the OAM Ethernet CFM maintenance domain MIP half functions.

NOTE: Whenever a MIP is configured and a bridge domain is mapped to multiple maintenance domains or maintenance associations, it is essential that the **mip-half-function** value for all maintenance domains and maintenance associations are the same.

Options **none**—Specify to not use the mip-half-function.

default—Specify to use the default mip-half-function.

explicit—Specify an explicit mip-half-function.

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

Related Documentation

- *Creating a Maintenance Domain*
- maintenance-domain on page 767

 Copyright © 2018, Juniper Networks, Inc.

mlfr-uni-nni-bundle-options

Syntax

```
mlfr-uni-nni-bundle-options {
  acknowledge-retries number;
  acknowledge-timer milliseconds;
  action-red-differential-delay (disable-tx | remove-link);
  drop-timeout milliseconds;
  fragment-threshold bytes;
  hello-timer milliseconds;
  link-layer-overhead percent;
  lmi-type (ansi | itu | c-lmi);
  minimum-links number;
  mrru bytes;
  n391 number;
  n392 number;
  n393 number;
  red-differential-delay milliseconds;
  t391 seconds;
  t392 number;
  yellow-differential-delay milliseconds;
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure link services and voice services interface management properties.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *Configuring Frame Relay Keepalives*
- lmi (Frame Relay) on page 738
- lmi-type on page 741
- *Junos OS Services Interfaces Library for Routing Devices*

mode (Dynamic Profiles)

Syntax mode loose;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family (inet) rpf-check]

Release Information Statement introduced in Junos OS Release 9.6.

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

- Configuring Unicast RPF Strict Mode on page 244

 Copyright © 2018, Juniper Networks, Inc.

mode (Interfaces)

Syntax mode loose;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family (inet | inet6) rpf-check],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family (inet | inet6) rpf-check]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 15.1F6 for PTX Series routers with third-generations FPCs installed.

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
- Configuring Unicast RPF Strict Mode on page 244

modem-options

Syntax
```
modem-options {
  dialin (console | routable);
  init-command-string initialization-command-string;
}
```

Hierarchy Level [edit interfaces umd0]

Release Information Statement introduced in Junos OS Release 8.2.

Description For J Series Services Routers, configure a USB port to act as a USB modem.

The remaining statement is explained separately. See CLI Explorer.

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

Related Documentation
- Specifying a USB Modem Interface on J Series Routers on page 350

Copyright © 2018, Juniper Networks, Inc.

monitor-session

Syntax monitor-session (*interface-name* | all);

Hierarchy Level [edit protocols ppp]

Release Information Statement introduced in Junos OS Release 7.5.

Description Monitor PPP packet exchanges. When monitoring is enabled, packets exchanged during a session are logged to the default log of **/var/log/pppd**.

Default If you do not include this statement, no PPPD-specific monitoring operations are performed.

Options **all**—Monitor PPP packet exchanges on all sessions.

interface-name—Logical interface name on which to enable session monitoring.

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

Related Documentation

- Monitoring a PPP Session on page 152

Copyright © 2018, Juniper Networks, Inc.

mpls (Interfaces)

Syntax

```
mpls {
  pop-all-labels {
    required-depth number;
  }
}
```

Hierarchy Level

[edit interfaces *interface-name* atm-options],
[edit interfaces *interface-name* sonet-options],
[edit interfaces *interface-name* fastether-options],
[edit interfaces *interface-name* gigether-options]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

For passive monitoring on ATM and SONET/SDH interfaces and 10-Gigabit Ethernet interfaces in WAN PHY mode, process incoming IP packets that have MPLS labels.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Removing MPLS Labels from Incoming Packets*
- *Enabling Packet Flow Monitoring on SONET/SDH Interfaces*
- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2018, Juniper Networks, Inc.

mrru

Syntax mrru *bytes*;

Hierarchy Level [edit interfaces *interface-name* mlfr-uni-nni-bundle-options],
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For multilink, link services, voice services, and J Series Services Routers ISDN interfaces only, set the maximum received reconstructed unit (MRRU). The MRRU is similar to the MTU, but is specific to multilink interfaces.

Options ***bytes***—MRRU size.

Range: 1500 through 4500 bytes

Default: 1500 bytes

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

Related Documentation

- mtu on page 785
- *Junos OS Services Interfaces Library for Routing Devices*

 Copyright © 2018, Juniper Networks, Inc.

mtu

Syntax mtu *bytes*;

Hierarchy Level [edit interfaces *interface-name*],
[edit interfaces *interface-name* unit *logical-unit-number* family *family*],
[edit interfaces interface-range *name*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*],
[edit logical-systems *logical-system-name* protocols l2circuit local-switching interface *interface-name* backup-neighbor *address*],
[edit logical-systems *logical-system-name* protocols l2circuit neighbor *address* interface *interface-name*],
[edit logical-systems *logical-system-name* protocols l2circuit neighbor *address* interface *interface-name* backup-neighbor *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols l2vpn interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls],
[edit protocols l2circuit local-switching interface *interface-name* backup-neighbor *address*],
[edit protocols l2circuit neighbor *address* interface *interface-name*]
[edit protocols l2circuit neighbor *address* interface *interface-name* backup-neighbor *address*],
[edit routing-instances *routing-instance-name* protocols l2vpn interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols vpls],
[edit logical-systems *name* protocols ospf area *name* interface],
[edit logical-systems *name* routing-instances *name* protocols ospf area *name* interface],
[edit protocols ospf area *name* interface],
[edit routing-instances *name* protocols ospf area *name* interface]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for Layer 2 VPNs and VPLS introduced in Junos OS Release 10.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.
Support at the **[set interfaces interface-name unit logical-unit-number family ccc]** hierarchy level introduced in Junos OS Release 12.3R3 for MX Series routers.
Statement introduced in Junos OS 17.3R1 Release for MX Series Routers.

Description Specify the maximum transmission unit (MTU) size for the media or protocol. The default MTU size depends on the device type. Changing the media MTU or protocol MTU causes an interface to be deleted and added again.

To route jumbo data packets on an integrated routing and bridging (IRB) interface or routed VLAN interface (RVI) on EX Series switches, you must configure the jumbo MTU size on the member physical interfaces of the VLAN that you have associated with the IRB interface or RVI, as well as on the IRB interface or RVI itself (the interface named irb or vlan, respectively).

Copyright © 2018, Juniper Networks, Inc.

CAUTION: For EX Series switches, setting or deleting the jumbo MTU size on an IRB interface or RVI while the switch is transmitting packets might cause packets to be dropped.

NOTE:

The MTU for an IRB interface is calculated by removing the Ethernet header overhead [6(DMAC)+6(SMAC)+2(EtherType)]. Because, the MTU is the lower value of the MTU configured on the IRB interface and the MTU configured on the IRB's associated bridge domain IFDs or IFLs, the IRB MTU is calculated as follows:

- In case of Layer 2 IFL configured with the **flexible-vlan-tagging** statement, the IRB MTU is calculated by including 8 bytes overhead (SVLAN+CVLAN).
- In case of Layer 2 IFL configured with the **vlan-tagging** statement, the IRB MTU is calculated by including a single VLAN 4 bytes overhead.

 Copyright © 2018, Juniper Networks, Inc.

NOTE:

- If a packet whose size is larger than the configured MTU size is received on the receiving interface, the packet is eventually dropped. The value considered for MRU (maximum receive unit) size is also the same as the MTU size configured on that interface.
- Not all devices allow you to set an MTU value, and some devices have restrictions on the range of allowable MTU values. You cannot configure an MTU for management Ethernet interfaces (fxp0, em0, or me0) or for loopback, multilink, and multicast tunnel devices.
- On ACX Series routers, you can configure the protocol MTU by including the **mtu** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet]** or **[edit interfaces *interface-name* unit *logical-unit-number* family inet6]** hierarchy level.
  - If you configure the protocol MTU at any of these hierarchy levels, the configured value is applied to all families that are configured on the logical interface.
  - If you are configuring the protocol MTU for both **inet** and **inet6** families on the same logical interface, you must configure the same value for both the families. It is not recommended to configure different MTU size values for **inet** and **inet6** families that are configured on the same logical interface.
- Starting in Release 14.2, MTU for IRB interfaces is calculated by removing the Ethernet header overhead **(6(DMAC)+6(SMAC)+2(EtherType))**, and the MTU is a minimum of the two values:
  - Configured MTU
  - Associated bridge domain's physical or logical interface MTU
    - For Layer 2 logical interfaces configured with **flexible-vlan-tagging**, IRB MTU is calculated by including 8 bytes overhead (**SVLAN+CVLAN**).
    - For Layer 2 logical interfaces configured with **vlan-tagging**, IRB MTU is calculated by including single VLAN 4 bytes overhead.

    NOTE: Changing the Layer 2 logical interface option from **vlan-tagging** to **flexible-vlan-tagging** or vice versa adjusts the logical interface MTU by 4 bytes with the existing MTU size. As a result, the Layer 2 logical interface is deleted and re-added, and the IRB MTU is re-computed appropriately.

For more information about configuring MTU for specific interfaces and router or switch combinations, see "Configuring the Media MTU" on page 115.

Options *bytes*—MTU size.

Range: 256 through 9192 bytes, 256 through 9216 (EX Series switch interfaces), 256 through 9500 bytes (Junos OS 12.1X48R2 for PTX Series routers), 256 through 9500 bytes (Junos OS 16.1R1 for MX Series routers)

NOTE: Starting in Junos OS Release 16.1R1, the MTU size for a media or protocol is increased from 9192 to 9500 for Ethernet interfaces on the following MX Series MPCs:

- MPC1
- MPC2
- MPC2E
- MPC3E
- MPC4E
- MPC5E
- MPC6E

Default: 1500 bytes (INET, INET6, and ISO families), 1448 bytes (MPLS), 1514 bytes (EX Series switch interfaces)

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

Related Documentation

- Configuring the Media MTU on page 115
- *Configuring the MTU for Layer 2 Interfaces*
- Setting the Protocol MTU on page 224

 Copyright © 2018, Juniper Networks, Inc.

multi-chassis-protection

Syntax

```
multi-chassis-protection {
  peer a.b.c.d {
    interface interface-name;
    }
  }
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 11.1.

Description For MX Series routers with multichassis aggregated Ethernet (MC-AE) interfaces, you can use this statement under the physical interface level to reduce the configuration at the logical interface level if the following assumption exists:

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

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

If the Inter-Chassis Control Protocol (ICCP) connection is UP and the interchassis data link (ICL) comes UP, the router configured as standby will bring up the MC-AE interfaces shared with the peer.

The remaining statements are explained separately. See CLI Explorer.

Options **interface *interface-name***—Specify the interface: **interface *interface-name-fpc/pic/port***

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

Related Documentation

- *Configuring Multichassis Link Aggregation on MX Series Routers*
- *Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation*
- *Configuring Aggregated Ethernet Link Protection*
- *Example: Configuring Aggregated Ethernet Link Protection*
- peer on page 864

multicast-dlci

Syntax multicast-dlci *dlci-identifier*;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description For point-to-multipoint Frame Relay, link services, and voice services interfaces only, enable multicast support on the interface. You can configure multicast support on the interface if the Frame Relay switch performs multicast replication.

Options ***dlci-identifier***—DLCI identifier, a number from 16 through 1022 that defines the Frame Relay DLCI over which the switch expects to receive multicast packets for replication.

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

Related Documentation

- *Configuring a Multicast-Capable Frame Relay Connection*
- dlci on page 540
- multipoint-destination on page 794
- *Junos OS Services Interfaces Library for Routing Devices*

 Copyright © 2018, Juniper Networks, Inc.

multicast-only

Syntax multicast-only;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the unit and family so that it can transmit and receive multicast traffic only. You can configure this property on the IP family only.

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

Related Documentation

- Configuring the Protocol Family on page 206
- *Junos OS Services Interfaces Library for Routing Devices*
- tunnel on page 1049

multicast-statistics

Syntax multicast-statistics;

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2.

Description For Ethernet, SONET, aggregated Ethernet, and aggregated SONET interfaces in T Series or TX Matrix routers, specify support for multicast statistics on a physical interface to enable multicast accounting for all the logical interfaces below the physical interface.

Default not enabled—must be configured to enable

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

Related Documentation

- *Configuring Multicast Statistics Collection on Aggregated Ethernet Interfaces*
- *Configuring Multicast Statistics Collection on Aggregated SONET Interfaces*
- *Configuring Multicast Statistics Collection on Ethernet Interfaces*
- *Configuring Multicast Statistics Collection on SONET Interfaces*

Copyright © 2018, Juniper Networks, Inc.

multicast-vci

Syntax multicast-vci *vpi-identifier*.*vci-identifier*;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description For ATM encapsulation only, and for point-to-multipoint ATM logical interfaces only, enable the support of multicast on the interface. You can configure multicast support on the interface if the ATM switch performs multicast replication.

Options ***vci-identifier***—ATM virtual circuit identifier.

Range: 0 through 16,384

vpi-identifier—ATM virtual path identifier.

Range: 0 through 255

Default: 0

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

Related Documentation

- *Configuring a Multicast-Capable ATM1 or ATM2 IQ Connection*
- multipoint-destination on page 794
- vci on page 1073

 Copyright © 2018, Juniper Networks, Inc.

multilink-max-classes

Syntax multilink-max-classes *number*;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description For Adaptive Services (AS) PIC link services IQ interfaces (**lsq**) only, configure the number of multilink classes to be negotiated when a link joins the bundle.

Options *number*—The number of multilink classes to be negotiated when a link joins the bundle.

Range: 1 through 8

Default: None

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*
- multipoint on page 793

multipoint

Syntax multipoint;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the interface unit as a multipoint connection.

Default If you omit this statement, the interface unit is configured as a point-to-point connection.

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

Related Documentation

- Configuring a Multipoint Connection on page 193
- point-to-point on page 873

Copyright © 2018, Juniper Networks, Inc.

multipoint-destination

Syntax

```
multipoint-destination address dlci dlci-identifier;
multipoint-destination address {
  epd-threshold cells;
  inverse-arp;
  oam-liveness {
    down-count cells;
    up-count cells;
  }
  oam-period (disable | seconds);
  shaping {
    (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate burst
      length);
    queue-length number;
  }
  vci vpi-identifier.vci-identifier;
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

For point-to-multipoint Frame Relay or ATM interfaces only, enable the support of multicast on the interface. You can configure multicast support on the interface if the Frame Relay or ATM switch performs multicast replication.

Options

address—Address of the remote side of the point-to-multipoint connection.

dlci-identifier—For Frame Relay interfaces, the data-link connection identifier.

Range: 0 through 0xFFFFFF (24 bits)

vci-identifier—For ATM interfaces, the virtual circuit identifier.

Range: 0 through 16,384

vpi-identifier—For ATM interfaces, the virtual path identifier.

Range: 0 through 255

Default: 0

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Configuring a Point-to-Point ATM1 or ATM2 IQ Connection*
- *Configuring a Point-to-Multipoint Frame Relay Connection*

Copyright © 2018, Juniper Networks, Inc.

- dlci on page 540
- encapsulation (Logical Interface) on page 573

multiservice-options

Syntax

```
multiservice-options {
  (syslog | no-syslog);
  (core-dump | no-core-dump);
  (dump-on-flow-control);
  flow-control-options {
    down-on-flow-control;
    dump-on-flow-control;
    reset-on-flow-control;
  }
}
```

Hierarchy Level [edit interfaces mo-*fpc/pic/port*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For monitoring services interfaces only, configure multiservice-specific interface properties.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- Configuring Multiservice Physical Interface Properties on page 170
- *Junos OS Services Interfaces Library for Routing Devices*
- passive-monitor-mode on page 857

n391

Syntax n391 *number*;

Hierarchy Level [edit interfaces *interface-name* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services and voice services interfaces only, set the Frame Relay full status polling interval.

Options *number*—Polling interval.

Range: 1 through 255

Default: 6

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*
- n392 on page 797
- n393 on page 798
- timeslots on page 1016
- t392 on page 1006

 Copyright © 2018, Juniper Networks, Inc.

n392

Syntax n392 *number*;

Hierarchy Level [edit interfaces *interface-name* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services and voices interfaces only, set the Frame Relay error threshold, in number of errors.

Options ***number***—Error threshold.

Range: 1 through 10

Default: 3

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*
- n391 on page 796
- n393 on page 798
- timeslots on page 1016
- t392 on page 1006

n393

Syntax n393 *number*;

Hierarchy Level [edit interfaces *interface-name* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services and voices interfaces only, set the Frame Relay monitored event count.

Options *number*—Number of event count.

Range: 1 through 10

Default: 4

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*
- n391 on page 796
- n392 on page 797
- timeslots on page 1016
- t392 on page 1006

 Copyright © 2018, Juniper Networks, Inc.

name-format

Syntax name-format (character-string | none | dns | mac+2oct);

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name*]

Release Information Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

Description Specify the format of the maintenance domain name.

Options **character-string**—The name is an ASCII character string.

none—The maintenance domain name is not used.

dns—The name is in domain name service (DNS) format. For example: www.juniper.net.

mac+2oct—Name is the MAC address plus a two-octet maintenance association identifier. For example: 08:00:22:33:44:55.100.

Default: character-string

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

Related Documentation

- *Creating a Maintenance Association*
- *Creating a Maintenance Domain*

Copyright © 2018, Juniper Networks, Inc.

native-vlan-id

Syntax

native-vlan-id *vlan-id*;

Hierarchy Level (QFX Series and EX4600)

For platforms without ELS:

[edit interfaces (QFX Series) *interface-name* unit 0 family ethernet-switching]

For platforms with ELS:

[edit interfaces (QFX Series) *interface-name*]

Hierarchy Level (ACX Series, EX Series, SRX Series, M Series, MX Series, and T Series)

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

Hierarchy Level (SRX Series)

[edit interfaces*interface-name*]

Release Information

Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 9.5 for SRX Series.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.

Description

Configure the VLAN identifier to associate with untagged packets received on the physical interface of a trunk mode interface for the following:

- QFX Series and EX4600
- M Series routers with Gigabit Ethernet IQ PICs with SFP and Gigabit Ethernet IQ2 PICs with SFP configured for 802.1Q flexible VLAN tagging
- MX Series routers with Gigabit Ethernet DPCs and MICs, Tri-Rate Ethernet DPCs and MICs, and 10-Gigabit Ethernet DPCs and MICs and MPCs configured for 802.1Q flexible VLAN tagging
- T4000 routers with 100-Gigabit Ethernet Type 5 PIC with CFP
- EX Series switches with Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces

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

 Copyright © 2018, Juniper Networks, Inc.

When the **native-vlan-id** statement is included with the **flexible-vlan-tagging** statement, untagged packets are accepted on the same mixed VLAN-tagged port and on the interfaces that are configured for Q-in-Q tunneling.

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

To configure the logical interface, include the **vlan-id** statement (matching the **native-vlan-id** statement on the physical interface) at the **[edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level**.

NOTE: **Starting in Junos OS Release 17.1R1, you can send untagged traffic without a native VLAN ID to the remote end of the network. To do this, remove the native VLAN ID from the untagged traffic configuration by setting the no-native-vlan-insert statement. If you do not configure this statement, the native VLAN ID is added to the untagged traffic.**

Default By default, the untagged packets are dropped. That is, if you do not configure the **native-vlan-id** option, the untagged packets are dropped.

Options *vlan-id*—Numeric identifier of the VLAN.

Range: 1 through 4094

number—VLAN ID number.

Range: (ACX Series routers, SRX Series devices and EX Series switches) 0 through 4094.

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

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit Ethernet Interfaces (J-Web Procedure)*
- *Understanding Bridging and VLANs on Switches*
- *Enabling VLAN Tagging*
- *Configuring Access Mode on a Logical Interface*
- *Configuring the Native VLAN Identifier on Switches With ELS Support (CLI Procedure)*
- *Understanding Interfaces*
- *Understanding Q-in-Q Tunneling and VLAN Translation*
- *no-native-vlan-insert*
- *Sending Untagged Traffic Without VLAN ID to Remote End*
- *show ethernet-switching interfaces*
- *show vlans*
- flexible-vlan-tagging on page 618
- *Junos OS Network Interfaces Configuration Guide*

ncp-max-conf-req

Syntax ncp-max-conf-req *number*

Hierarchy Level [edit interfaces so-*fpc/pic/port* unit *number* ppp-options]

Release Information Statement introduced in Junos OS Release 9.6.

Description Set the maximum number of NCP Configure-Requests to be sent, after which the router goes to NCP down state.

Options ***number***—Ranges from 0 to 65535, where 0 means send infinite NCP Configure-Requests and any other value specifies the maximum number NCP Configure-Requests to send and then stop sending.

Default—254

Range: 0 through 65,535

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

Related Documentation

- Configuring the NCP Configure-Request Maximum Sent on page 195
- ppp-options on page 891

 Copyright © 2018, Juniper Networks, Inc.

ncp-restart-timer

Syntax ncp-restart-timer *milliseconds*;

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

Release Information Statement introduced in Junos OS Release 8.1.

Description For interfaces with PPP and PPP TCC encapsulations and on multilink PPP bundle interfaces, configure a restart timer for the Network Control Protocol (NCP) component of a PPP session.

Options ***milliseconds***—The time in milliseconds between successive NCP configuration requests.

Range: 500 through 10,000 milliseconds

Default: 3 seconds

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

Related Documentation

- Configuring the PPP Restart Timers on page 193

Copyright © 2018, Juniper Networks, Inc.

nd6-stale-time

Syntax nd6-stale-time *seconds*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet6]

Release Information Statement introduced in Junos OS Release 11.1.
Statement introduced in Junos OS Release 11.1 for EX Series switches.

Description Set the stale timer for IPv6 neighbor reachability confirmation. Reachability of the IPv6 neighbors is confirmed only after the stale timer has expired. For example, by setting the stale timer to 180 seconds, users can specify that IPv6 neighbor reachability be confirmed every 180 seconds.

NOTE: When the Routing Engine sends a control packet to an IPv6 neighbor, the stale timer is the maximum interval in which neighbor reachability is confirmed. In such cases, IPv6 neighbor reachability is confirmed before the stale timer expires.

Default Default is 20 minutes (1200 seconds)

Options *seconds*—Duration in seconds.

Range: 1 to 18000

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

Related Documentation

- *IPv6 Neighbor Discovery Overview*
- show ipv6 neighbors on page 1966

Copyright © 2018, Juniper Networks, Inc.

negotiate-address

Syntax negotiate-address;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description For interfaces with PPP encapsulation, enable the interface to be assigned an IP address by the remote end.

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

Related Documentation

- Configuring IPCP Options for Interfaces with PPP Encapsulation on page 216
- address on page 424
- unnumbered-address (PPP) on page 1067
- *Junos OS Administration Library*

negotiation-options

Syntax

```
negotiation-options {
  allow-remote-loopback;
  no-allow-link-events;
}
```

Hierarchy Level [edit protocols oam link-fault-management interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.4.

Description Enable and disable IEEE 802.3ah Operation, Administration, and Management (OAM) features for Ethernet interfaces.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *IEEE 802.3ah OAM Link-Fault Management Overview*

Copyright © 2018, Juniper Networks, Inc.

neighbor (Automatic Protection Switching for SONET/SDH)

Syntax neighbor *address*;

Hierarchy Level [edit interfaces *interface-name* sonet-options aps]

Release Information Statement introduced before Junos OS Release 7.4.

Description If you are configuring one router to be the working router and a second to be the protect router, configure the address of the remote interface. You configure this on one or both of the interfaces.

The address you specify for the neighbor must never be routed through the interface on which APS is configured, or instability will result. We strongly recommend that you directly connect the working and protect routers and that you configure the interface address of this shared network as the neighbor address.

Options ***address***—Neighbor's address.

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

Related Documentation

- *Configuring Basic Automatic Protect Switching*

no-allow-link-events

Syntax no-allow-link-events;

Hierarchy Level [edit protocols oam ethernet link-fault-management interface *interface-name* negotiation-options]

Release Information Statement introduced in Junos OS Release 8.4.

Description Disable the sending of link event TLVs.

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

Related Documentation

- *Disabling the Sending of Link Event TLVs*

 Copyright © 2018, Juniper Networks, Inc.

no-aggregate-delegate-processing

Syntax no-aggregate-delegate-processing;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management]

Release Information Statement introduced in Junos OS Release 14.1.

Description Disable distribution of connectivity fault management (CFM) sessions on aggregated Ethernet interfaces.

Default CFM sessions on aggregated Ethernet interfaces are distributed by default.

Required Privilege Level Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- *IEEE 802.1ag OAM Connectivity Fault Management Overview*

Copyright © 2018, Juniper Networks, Inc.

asynchronous-notification

Syntax (asynchronous-notification | no-asynchronous-notification);

Hierarchy Level [edit interfaces ge-*fpc/pic/port* gigether-options]

Release Information Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description (MX Series routers, T Series routers) For all Gigabit Ethernet interfaces (1-Gigabit, 10-Gigabit, and 100-Gigabit), configure support for notification of link down alarm generation and transfer.

(M120 and M320 routers) For all 10-Gigabit Ethernet PIC interfaces, configure support for notification of link down alarm generation and transfer.

- **asynchronous-notification**—Support notification of link down alarm generation and transfer.
- **no-asynchronous-notification**—Prohibit notification of link down alarm generation and transfer.

Default Support for notification of link down alarm generation and transfer is not enabled.

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

Related Documentation

- *Gigabit Ethernet Notification of Link Down Alarm Overview*
- *Configuring Gigabit Ethernet Notification of Link Down Alarm*

 Copyright © 2018, Juniper Networks, Inc.

no-auto-mdix

Syntax no-auto-mdix;

Hierarchy Level [edit interface ge-*fpc*/*port*/*pic* gigether-options]

Release Information Statement introduced in Junos OS Release 9.5.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description Disable the Auto MDI/MDIX feature.

MX Series routers with Gigabit Ethernet interfaces automatically detect MDI and MDIX port connections. Use this statement to override the default setting. Remove this statement to return to the default setting.

Default Auto MDI/MDIX is enabled by default.

Options There are no options for this statement.

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

Related Documentation

- *Ethernet Interfaces Overview*
- gigether-options on page 632.

Copyright © 2018, Juniper Networks, Inc.

auto-negotiation

Syntax

(auto-negotiation | no-auto-negotiation) <remote-fault (local-interface-online | local-interface-offline)>;

Hierarchy Level

[edit interfaces *interface-name* ether-options],
[edit interfaces *interface-name* gigether-options],
[edit interfaces ge-*pim*/0/0 switch-options switch-port *port-number*]

Release Information

Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description

For Gigabit Ethernet interfaces on M Series, MX Series, T Series, TX Matrix routers, and ACX Series routers explicitly enable autonegotiation and remote fault. For EX Series switches, explicitly enable autonegotiation only.

- **auto-negotiation**—Enables autonegotiation. This is the default.
- **no-auto-negotiation**—Disable autonegotiation. When autonegotiation is disabled, you must explicitly configure the link mode and speed.

When you configure Tri-Rate Ethernet copper interfaces to operate at 1 Gbps, autonegotiation must be enabled.

NOTE: On EX Series switches, an interface configuration that disables autonegotiation and manually sets the link speed to 1 Gbps is accepted when you commit the configuration; however, if the interface you are configuring is a Tri-Rate Ethernet copper interface, the configuration is ignored as invalid and autonegotiation is enabled by default.

To correct the invalid configuration and disable autonegotiation:

1. Delete the **no-auto-negotiation** statement and commit the configuration.
2. Set the link speed to 10 or 100 Mbps, set **no-auto-negotiation**, and commit the configuration.

On EX Series switches, if the link speed and duplex mode are also configured, the interfaces use the values configured as the desired values in the negotiation. If autonegotiation is disabled, the link speed and link mode must be configured.

NOTE: On T4000 routers, the **auto-negotiation** command is ignored for interfaces other than Gigabit Ethernet.

 Copyright © 2018, Juniper Networks, Inc.

NOTE: On ACX Series routers, when you configure fiber interfaces (fiber media mode) to operate at 1 Gbps, you need to always enable autonegotiation (**auto-negotiation**) to negotiate the speed and duplex settings. You cannot disable autonegotiation (**no-auto-negotiation**) in the fiber media mode. In copper interfaces (copper media mode), autonegotiation is enabled by default. To disable autonegotiation, you need to explicitly configure the link speed to 10 or 100 Mbps, set **no-auto-negotiation**, and commit the configuration.

Default Autonegotiation is automatically enabled. No explicit action is taken after the autonegotiation is complete or if the negotiation fails.

Options **remote-fault (local-interface-online | local-interface-offline)**—(Optional) For M Series, MX Series, T Series, TX Matrix routers, and ACX Series routers only, manually configure remote fault on an interface.

Default: local-interface-online

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

Related Documentation

- *Gigabit Ethernet Autonegotiation Overview*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure) for EX Series Switches with ELS support*

Copyright © 2018, Juniper Networks, Inc.

cbit-parity

Syntax (cbit-parity | no-cbit-parity);

Hierarchy Level [edit interfaces *interface-name* t3-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For T3 interfaces only, enable or disable C-bit parity mode, which controls the type of framing that is present on the transmitted T3 signal. When C-bit parity mode is enabled, the C-bit positions are used for the far-end block error (FEBE), far-end alarm and control (FEAC), terminal data link, path parity, and mode indicator bits, as defined in ANSI T1.107a-1989. For ATM and ATM2 IQ2 and IQ2-E interfaces, M23 framing is used when the **no-cbit-parity** statement is included. For all other interfaces, M13 framing is used when the **no-cbit-parity** statement is included.

Default C-bit parity mode is enabled.

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

Related Documentation

- *Configuring E3 and T3 Parameters on ATM Interfaces*
- *Disabling T3 C-Bit Parity Mode*

Copyright © 2018, Juniper Networks, Inc.

core-dump

Syntax (core-dump | no-core-dump);

Hierarchy Level [edit interfaces mo-*fpc/pic/port* multiservice-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For monitoring services interfaces only, a useful tool for isolating the cause of a problem. Core dumping is enabled by default. The directory **/var/tmp** contains core files. The Junos OS saves the current core file (0) and the four previous core files, which are numbered 1 through 4 (from newest to oldest):

- **core-dump**—Enable the core dumping operation.
- **no-core-dump**—Disable the core dumping operation.

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

Related Documentation

- Configuring Multiservice Physical Interface Properties on page 170
- *Junos OS Services Interfaces Library for Routing Devices*

feac-loop-respond

Syntax (feac-loop-respond | no-feac-loop-respond);

Hierarchy Level [edit interfaces *interface-name* t3-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For T3 interfaces only, configure the router so a remote CSU can place the local router into loopback.

If you configure remote or local loopback with the T3 **loopback** statement, the router does not respond to FEAC requests from the CSU even if you include the **feac-loop-respond** statement in the configuration. For the router to respond, you must delete the **loopback** statement from the configuration.

You must rollback the setting done on the remote CSU prior to deactivating the **feac-loop-respond** statement. If the remote CSU cannot comply, clear the remote loop through local configuration to achieve the cleanup. For example, configure remote loopback on the interface and then delete the remote loopback.

Default The router does not respond to FEAC requests.

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

Related Documentation

- *Configuring the T3 FEAC Response*
- loopback (ADSL, DS0, E1/E3, SONET/SDH, SHDSL, and T1/T3) on page 751
- remote-loopback-respond on page 928

 Copyright © 2018, Juniper Networks, Inc.

flow-control

Syntax (flow-control | no-flow-control);

Hierarchy Level [edit interfaces *interface-name* aggregated-ether-options],
[edit interfaces *interface-name* ether-options],
[edit interfaces *interface-name* fastether-options],
[edit interfaces *interface-name* gigether-options],
[edit interfaces *interface-name* multiservice-options],
[edit interfaces interface-range *name* aggregated-ether-options],
[edit interfaces interface-range *name* ether-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 in EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description For aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only, explicitly enable flow control, which regulates the flow of packets from the router or switch to the remote side of the connection. Enabling flow control is useful when the remote device is a Gigabit Ethernet switch. Flow control is not supported on the 4-port Fast Ethernet PIC.

NOTE: On the Type 5 FPC, to prioritize control packets in case of ingress oversubscription, you must ensure that the neighboring peers support MAC flow control. If the peers do not support MAC flow control, then you must disable flow control.

Default Flow control is enabled.

NOTE: Flow control is enabled by default only on physical interfaces and it is disabled by default on aggregated Ethernet interfaces.

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

Related Documentation

- *Configuring Flow Control*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure) for EX Series Switches with ELS support*

Copyright © 2018, Juniper Networks, Inc.

gratuitous-arp-reply

Syntax (gratuitous-arp-reply | no-gratuitous-arp-reply);

Hierarchy Level [edit interfaces *interface-name*]
[edit interfaces **interface-range** *interface-range-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 in EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description For Ethernet interfaces, enable updating of the Address Resolution Protocol (ARP) cache for gratuitous ARPs.

Default Updating of the ARP cache is disabled on all Ethernet interfaces.

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

Related Documentation

- *Configuring Gratuitous ARP*
- no-gratuitous-arp-request on page 816

no-gratuitous-arp-request

Syntax no-gratuitous-arp-request;

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 9.6 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description For Ethernet interfaces and pseudowire logical interfaces, do not respond to gratuitous ARP requests.

Default Gratuitous ARP responses are enabled on all Ethernet interfaces.

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

Related Documentation

- *Configuring Gratuitous ARP*

Copyright © 2018, Juniper Networks, Inc.

no-keepalives

Syntax

no-keepalives;

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Disable the sending of keepalives on a physical interface configured with PPP, Frame Relay, or Cisco HDLC encapsulation. The default keepalive interval is 10 seconds.

For ATM2 IQ interfaces only, you can disable keepalives on a logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:

- **atm-ppp-llc**—PPP over AAL5 LLC encapsulation.
- **atm-ppp-vc-mux**—PPP over AAL5 multiplex encapsulation.

Required Privilege Level

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

Related Documentation

- Configuring Keepalives on page 134
- *Disabling the Sending of PPPoE Keepalive Messages*
- *Configuring Frame Relay Keepalives*

Copyright © 2018, Juniper Networks, Inc.

long-buildout

Syntax (long-buildout | no-long-buildout);

Hierarchy Level [edit interfaces *interface-name* t3-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the T3 line buildout. A T3 interface has two settings for the T3 line buildout: a short setting, which is less than 255 feet (68 meters), and a long setting, which is greater than 255 feet and shorter than 450 feet (137 meters).

This statement applies to copper-cable-based T3 interfaces only. You cannot configure a line buildout for a DS3 channel on a channelized OC12 interface, which runs over fiber-optic cable.

Default A T3 interface uses the short line buildout setting (**no-long-buildout**) for wires shorter than 255 feet (68 meters).

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

Related Documentation

- *Configuring the T3 Line Buildout*

 Copyright © 2018, Juniper Networks, Inc.

loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet)

Syntax (loopback | no-loopback);

Hierarchy Level [edit interfaces *interface-name* aggregated-ether-options],
[edit interfaces *interface-name* ether-options],
[edit interfaces *interface-name* fastether-options],
[edit interfaces *interface-name* gigether-options],
[edit interfaces interface-range *name* ether-options]

For QFX Series and EX Series:

[edit interfaces *interface-name* aggregated-ether-options],
[edit interfaces *interface-name* ether-options],

For SRX Series Devices and vSRX:

[edit interfaces *interface-name* redundant-ether-options]

Release Information Statement introduced before Junos OS Release 7.4 for MX Series.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Statement modified in Junos OS Release 9.2 for the SRX Series.

Description For aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces, enable or disable loopback mode.

NOTE:

- By default, local aggregated Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces connect to a remote system.
- IPv6 Neighbor Discovery Protocol (NDP) addresses are not supported on Gigabit Ethernet interfaces when loopback mode is enabled on the interface. That is, if the **loopback** statement is configured at the **[edit interfaces ge-*fpc/pic/port* gigether-options]** hierarchy level, an NDP address cannot be configured at the **[edit interfaces ge-*fpc/pic/port* unit *logical-unit-number* family inet6 address]** hierarchy level.

Default By default, loopback is disabled.

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

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- *Configuring Ethernet Loopback Capability*
- *Understanding Interfaces*

mac-learn-enable

Syntax

mac-learn-enable;

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile]
[edit interfaces ae*x* aggregated-ether-options ethernet-switch-profile]

Release Information

Statement introduced before Junos OS Release 7.4.
Support for statement under the **[edit interfaces ae*x* aggregated-ether-options ethernet-switch-profile]** hierarchy introduced in Junos OS Release 15.1.

Description

For Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), for Gigabit Ethernet DPCs on MX Series routers, for 100-Gigabit Ethernet Type 5 PIC with CFP, and for MPC3E, MPC4E, MPC5E, MPC5EQ, and MPC6E MPCs, configure dynamic learning of the source and destination MAC addresses. By default, the interface is not allowed to dynamically learn source and destination MAC addresses.

To disable dynamic learning of the source and destination MAC addresses after it has been configured, you must delete **mac-learn-enable** from the configuration.

MPCs support MAC address accounting for an individual interface or an aggregated Ethernet interface member link only after the interface has received traffic from the MAC source. If traffic is only exiting an interface, the MAC address is not learned and MAC address accounting does not occur.

Required Privilege Level

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

Related Documentation

- *Configuring MAC Address Filtering*
- *Configuring MAC Address Accounting*

 Copyright © 2018, Juniper Networks, Inc.

no-partition

Syntax

no-partition interface-type (e1 | (cau4 | so) | (ct3 | t3) | so | t3);

Hierarchy Level

[edit interfaces ce1-*fpc/pic/port*],
[edit interfaces coc1-*fpc/pic/port*:*channel*],
[edit interfaces coc12-*fpc/pic/port*],
[edit interfaces cstm1-*fpc/pic/port*],
[edit interfaces ct3-*fpc/pic/port*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

For Channelized E1 IQ PICs only, configure the channelized E1 interface as an unpartitioned, clear channel.

For Channelized OC12 PIC only, convert the channelized OC1 IQ interface into a channelized T3 interface or a T3 interface. You perform this configuration task for C-bit parity and M13-mapped configurations.

For Channelized OC12 IQ PICs only, configure the channelized OC12 interface as an unpartitioned, clear channel.

For Channelized STM1 PIC only, convert the channelized STM1 IQ interface into a channelized Administrative Unit 4 (AU-4) interface or a SONET/SDH STM1 interface.

For Channelized DS3 PIC only, configure the channelized T3 interface as an unpartitioned, clear channel.

Default

If you do not include either this statement or the **partition** statement, the Channelized IQ PIC is not partitioned, and no data channels are configured.

Options

The option used must correspond to the physical interface type:

e1—E1 interface type.

coc12 so—Channelized OC12 interface type, in SONET mode.

cau4—Channelized AU-4 interface type.

cstm1—SONET/SDH STM1 interface type, in SDH mode.

ct3—Channelized T3 interface type.

t3—T3 interface type.

Required Privilege Level

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

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- *Channelized E1 IQ and IQE Interfaces Overview*
- *Channelized OC12/STM4 IQ and IQE Interfaces Overview*
- *Configuring an OC12/STM4 Interface*
- *Configuring Channelized STM1 IQ and IQE Interfaces*
- *Configuring T3 IQ Interfaces*
- partition on page 854
- *no-partition*

Copyright © 2018, Juniper Networks, Inc.

payload-scrambler

Syntax (payload-scrambler | no-payload-scrambler);

Hierarchy Level [edit interfaces *interface-name* e3-options],
[edit interfaces *interface-name* sonet-options],
[edit interfaces *interface-name* t3-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Enable or disable HDLC scrambling on an E3, a SONET/SDH, or a T3 interface. This type of scrambling provides better link stability. Both sides of a connection must either use or not use scrambling.

If you commit a T3 interface configuration that has HDLC payload scrambling enabled, the interface must also be configured to be compatible with the channel service unit (CSU) at the remote end of the line.

Disable payload scrambling on an E3 interface if Digital Link compatibility mode is used.

On a channelized OC12 interface, the **sonet payload-scrambler** statement is ignored. To configure scrambling on the DS3 channels on the interface, you can include the **t3-options payload-scrambler** statement in the configuration for each DS3 channel.

NOTE: The **payload-scrambler** statement at the **[edit interfaces *interface-name* e3-options]** hierarchy level is not valid for IQE PICs.

Default Payload scrambling is disabled on all E3 and T3 interfaces; it is enabled by default on E3/T3 over ATM interfaces and on SONET/SDH interfaces.

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

Related Documentation

- *Configuring E3 and T3 Parameters on ATM Interfaces*
- *Configuring E3 HDLC Payload Scrambling*
- *Configuring SONET/SDH HDLC Payload Scrambling for Link Stability*
- *Configuring T3 HDLC Payload Scrambling*
- *Examples: Configuring T3 Interfaces*
- compatibility-mode on page 491

Copyright © 2018, Juniper Networks, Inc.

no-pre-classifier

Syntax no-pre-classifier;

Hierarchy Level [edit chassis fpc *n* pic *n*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify disabling the control queue for all ports on the 10-Gigabit Ethernet LAN/WAN PIC. Deleting this configuration re-enables the control queue feature on all ports of the 10-Gigabit Ethernet LAN/WAN PIC.

NOTE: For the 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PD-5-10XGE-SFPP), the control queue has a rate limiter to limit the control traffic to 2 Mbps (fixed, not user-configurable) per port. If the transit control traffic crosses this limit, then it can cause drops on locally terminating control traffic, causing flap of protocols such as BGP and OSPF. To avoid the control traffic being dropped, configure the **no-pre-classifier** statement to disable the control queue.

Default The **no-pre-classifier** statement is not configured and the control queue is operational.

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

Related Documentation

- *10-port 10-Gigabit Ethernet LAN/WAN PIC Overview*
- *Configuring Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC*

 Copyright © 2018, Juniper Networks, Inc.

no-redirects

Syntax no-redirects;

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

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Do not send protocol redirect messages on the interface.

To disable the sending of protocol redirect messages for the entire router or switch, include the **no-redirects** statement at the **[edit system]** hierarchy level.

Default Interfaces send protocol redirect messages.

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

Related Documentation

- Disabling the Transmission of Redirect Messages on an Interface on page 226
- *Junos OS Administration Library*

Copyright © 2018, Juniper Networks, Inc.

source-filtering

Syntax

(source-filtering | no-source-filtering);

Hierarchy Level

[edit interfaces *interface-name* aggregated-ether-options],
[edit interfaces *interface-name* fastether-options],
[edit interfaces *interface-name* gigether-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Packet Transport Routers.

Description

For aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, and Gigabit Ethernet IQ interfaces only, enable the filtering of MAC source addresses, which blocks all incoming packets to that interface. To allow the interface to receive packets from specific MAC addresses, include the **source-address-filter** statement.

If the remote Ethernet card is changed, the interface is no longer able to receive packets from the new card because it has a different MAC address.

Default

Source address filtering is disabled.

Required Privilege Level

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

Related Documentation

- *Configuring MAC Address Filtering for Ethernet Interfaces*
- *Configuring MAC Address Filtering on PTX Series Packet Transport Routers*
- accept-source-mac on page 409
- source-address-filter on page 976

 Copyright © 2018, Juniper Networks, Inc.

syslog (Monitoring)

Syntax (syslog | no-syslog);

Hierarchy Level [edit interfaces mo-*fpc/pic/port* multiservice-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description System logging is enabled by default. The system log information of the Monitoring Services PIC is passed to the kernel for logging in the **/var/log** directory.

- **syslog**—Enable PIC system logging.
- **no-syslog**—Disable PIC system logging.

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

Related Documentation

- Configuring Multiservice Physical Interface Properties on page 170
- *Junos OS Services Interfaces Library for Routing Devices*

no-termination-request

Syntax no-termination-request;

Hierarchy Level [edit interfaces *interface-name* ppp-options],
[edit interfaces lsq-*fpc/pic/port* lsq-failure-options]

Release Information Statement introduced in Junos OS Release 7.4.
Support at the **[edit interfaces *interface-name* ppp-options]** hierarchy level added in Junos OS Release 8.3.

Description For LSQ PICs or link PICs in redundant LSQ configurations, you can inhibit the router from sending PPP termination-request messages to the remote host if the PIC fails.

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

Related Documentation

- *Configuring Link PIC Failover on Channelized OC3 IQ and IQE Interfaces*
- *Configuring Link PIC Failover on Channelized OC12/STM4 IQ and IQE Interfaces*
- *Configuring Link PIC Failover on Channelized STM1 Interfaces*
- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2018, Juniper Networks, Inc.

translate-discard-eligible

Syntax (translate-discard-eligible | no-translate-discard-eligible);

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

Release Information Statement introduced before Junos OS Release 7.4.

Description For interfaces with encapsulation type Frame Relay CCC, enable or disable translation of Frame Relay discard eligible (DE) control bits.

Default DE bit translation is disabled.

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

Related Documentation
- *Configuring Frame Relay Control Bit Translation*

translate-fecn-and-becn

Syntax (translate-fecn-and-becn | no-translate-fecn-and-becn);

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

Release Information Statement introduced before Junos OS Release 7.4.

Description For interfaces with encapsulation type Frame Relay CCC, enable or disable translation of Frame Relay forward explicit congestion notification (FECN) control bits and Frame Relay backward explicit congestion notification (BECN) control bits.

Default FECN and BECN bit translation is disabled.

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

Related Documentation
- *Configuring Frame Relay Control Bit Translation*

 Copyright © 2018, Juniper Networks, Inc.

unframed

Syntax (unframed | no-unframed);

Hierarchy Level [edit interfaces *interface-name* e3-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For E3 IQ interfaces only, enable or disable unframed mode. In unframed mode, the E3 IQ interface do not detect yellow (**ylw**) or loss-of-frame (**lof**) alarms.

Default Unframed mode is disabled.

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

Related Documentation
- *Configuring E3 IQ and IQE Unframed Mode*

z0-increment

Syntax (z0-increment | no-z0-increment);

Hierarchy Level [edit interfaces *interface-name* sonet-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure an incremental STM ID rather than a static one.

Default **no-Z0-increment**

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

Related Documentation
- *Configuring an Incrementing STM ID to Interoperate with Older Equipment in SDH Mode*
- sonet-options on page 973

node-id

Syntax node-id *mac-address*;

Hierarchy Level [edit protocols protection-group ethernet-ring *ring-name*]

Release Information Statement introduced in Junos OS Release 9.4.

Description For EX Series switches and QFX Series switches, node-id is not configurable.

For MX Series routers, optionally specify the MAC address of a node in the protection group. If this statement is not included, the router assigns the node's MAC address.

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

Related Documentation

- *Ethernet Ring Protection Switching Overview*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*

non-revertive (Interfaces)

Syntax non-revertive;

Hierarchy Level [edit interfaces ae*X* aggregated-ether-options lacp link-protection]

Release Information Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 11.4 for EX Series switches.
Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.

Description Disable the ability to switch to a better priority link (if one is available) once a link is established as active and collection distribution is enabled.

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

Related Documentation

- link-protection on page 733
- *Configuring Aggregated Ethernet Link Protection*
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*

Copyright © 2018, Juniper Networks, Inc.

oam

Syntax

```
oam {
  ethernet {
    connectivity-fault-management {
      action-profile profile-name {
        default-actions {
          interface-down;
        }
      }
      performance-monitoring {
        delegate-server-processing;
        hardware-assisted-timestamping;
        hardware-assisted-keepalives;
        sla-iterator-profiles {
          profile-name {
            avg-fd-twoway-threshold;
            avg-ifdv-twoway-threshold;
            avg-flr-forward-threshold;
            avg-flr-backward-threshold;
            disable;
            calculation-weight {
              delay delay-weight;
              delay-variation delay-variation-weight;
            }
            cycle-time milliseconds;
            iteration-period connections;
            measurement-type (loss | statistical-frame-loss | two-way-delay);
          }
        }
      }
      linktrace {
        age (30m | 10m | 1m | 30s | 10s);
        path-database-size path-database-size;
      }
      maintenance-domain domain-name {
        level number;
        name-format (character-string | none | dns | mac+2octet);
        maintenance-association ma-name {
          short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
          protect-maintenance-association protect-ma-name;
          remote-maintenance-association remote-ma-name;
          continuity-check {
            convey-loss-threshold;
            hold-interval minutes;
            interface-status-tlv;
            interval (100ms | 10m | 10ms | 10s | 1m | 1s);
            loss-threshold number;
            port-status-tlv;
          }
          mep mep-id {
            auto-discovery;
            direction (up | down);
            interface interface-name (protect | working);
```

```
                lowest-priority-defect (all-defects | err-xcon | mac-rem-err-xcon | no-defect |
                  rem-err-xcon | xcon );
                priority number;
                remote-mep mep-id {
                  action-profile profile-name;
                  sla-iterator-profile profile-name {
                    data-tlv-size size;
                    iteration-count count-value;
                    priority priority-value;
                  }
                }
              }
            }
          }
        }
        link-fault-management {
          action-profile profile-name {
            action {
              link-down;
              send-critical-event;
              syslog;
            }
            event {
              link-adjacency-loss;
              link-event-rate {
                frame-error count;
                frame-period count;
                frame-period-summary count;
                symbol-period count;
              }
              protocol-down;
            }
          }
          interface interface-name {
            apply-action-profile
            link-discovery (active | passive);
            loopback-tracking;
            pdu-interval interval;
            pdu-threshold threshold-value;
            remote-loopback;
            event-thresholds {
              frame-error count;
              frame-period count;
              frame-period-summary count;
              symbol-period count;
            }
            negotiation-options {
              allow-remote-loopback;
              no-allow-link-events;
            }
          }
        }
      }
    }
```

 Copyright © 2018, Juniper Networks, Inc.

Hierarchy Level [edit protocols]

Release Information Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

Description For Ethernet interfaces on M320, M120, MX Series, and T Series routers and PTX Series Packet Transport Routers, provide IEEE 802.3ah Operation, Administration, and Maitenance (OAM) support.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *IEEE 802.3ah OAM Link-Fault Management Overview*

oam-liveness

Syntax

```
oam-liveness {
  down-count cells;
  up-count cells;
}
```

Hierarchy Level

[edit interfaces *interface-name* atm-options vpi *vpi-identifier*],
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit interfaces *interface-name* unit *logical-unit-number* family *family* address *address* multipoint-destination *address*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family* address *address* multipoint-destination *address*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description

For ATM encapsulation only, configure Operation, Administration, and Maintenance (OAM) F5 loopback cell count thresholds. Not supported on ATM-over-SHDSL interfaces.

For ATM2 IQ PICs only, configure OAM F4 loopback cell count thresholds at the **[edit interfaces *interface-name* atm-options vpi *vpi-identifier*]** hierarchy level.

Options

down-count *cells*—Minimum number of consecutive OAM F4 or F5 loopback cells lost before a VC is declared down.

Range: 1 through 255

Default: 5 cells

up-count *cells*—Minimum number of consecutive OAM F4 or F5 loopback cells received before a VC is declared up.

Range: 1 through 255

Default: 5 cells

Required Privilege Level

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

Related Documentation

- *Configuring the ATM OAM F5 Loopback Cell Threshold*

Copyright © 2018, Juniper Networks, Inc.

oam-period

Syntax oam-period (disable | *seconds*);

Hierarchy Level [edit interfaces *interface-name* atm-options vpi *vpi-identifier*],
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit interfaces *interface-name* unit *logical-unit-number* family *family* address *address*
multipoint-destination *address*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*
family *family* address *address* multipoint-destination *address*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description For ATM encapsulation only, configure the OAM F5 loopback cell period. Not supported on ATM-over-SHDSL interfaces.

For ATM2 IQ PICs only, configure the OAM F4 loopback cell period at the **[edit interfaces *interface-name* atm-options vpi *vpi-identifier*]** hierarchy level.

Default If you omit this statement, OAM F5 loopback cells are not initiated, but the interface still responds if it receives OAM F5 loopback cells.

Options **disable**—Disable the OAM loopback cell transmit feature.

seconds—OAM loopback cell period.

Range: 1 through 900 seconds

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

Related Documentation
- *Defining the ATM OAM F5 Loopback Cell Period*

Copyright © 2018, Juniper Networks, Inc.

oc-slice

Syntax oc-slice *oc-slice-range*;

Hierarchy Level [edit interfaces *interface-name* partition *partition-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For channelized OC12 IQ interfaces only, configure the range of SONET/SDH slices.

Default If you do not include either this statement or the **no-partition** statement, the Channelized OC12 IQ PICs not partitioned, and no data channels are configured.

Options ***oc-slice-range***—Range of SONET/SDH slices. OC3 interfaces must occupy three consecutive OC slices per interface, in the form 1–3, 4–6, 7–9, or 10–12. The T3, T1, and DS0 interface types each occupy one OC slice per interface.

Range: For OC3 interfaces,1–3, 4–6, 7–9, or 10–12; for SONET/SDH and T3 interfaces, 1–12

The remaining statement is explained separately. See CLI Explorer.

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

Related Documentation
- *Channelized OC12/STM4 IQ and IQE Interfaces Overview*

 Copyright © 2018, Juniper Networks, Inc.

open-timeout

Syntax open-timeout *seconds*;

Hierarchy Level [edit interfaces *interface-name* services-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure timeout period for Transmission Control Protocol (TCP) session establishment.

Options *seconds*—Timeout period in seconds.

Range: 4 through 224 seconds

Default: 5 seconds

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

operating-mode

Syntax operating-mode *mode*;

Hierarchy Level [edit interfaces at-*fpc/pic/port* dsl-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For J Series Services Routers only, modify the operating mode of the digital subscriber line for an ATM interface.

Options ***mode***—Operating mode for ATM-over-ADSL interfaces. The mode can be one of the following:

- **adsl2plus**—Set the ADSL line to train in the ITU G.992.5 mode.
- **ansi-dmt**—Set the ADSL line to train in the ANSI T1.413 Issue 2 mode.
- **auto**—Set the ADSL line to autonegotiate the setting to match the setting of the DSL access multiplexer (DSLAM) located at the central office. The ADSL line trains in the ANSI T1.413 Issue 2 (**ansi-dmt**) or ITU G.992.1 (**itu-dmt**) mode.
- **etsi**—Set the ADSL line to train in the ETSI TS 101 388 V1.3.1 mode.
- **itu-annexb-ur2**—Set the ADSL line to train in the ITU G.992.1 UR-2 mode.
- **itu-annexb-non-ur2**—Set the ADSL line to train in the ITU G.992.1 non-UR-2 mode.
- **itu-dmt**—Set the ADSL line to train in the ITU G.992.1 mode.
- **itu-dmt-bis**—Set the ADSL line to train in the ITU G.992.3 mode.

Default: auto

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

Related Documentation

- *ATM-over-ADSL Overview*
- *Junos OS Interfaces and Routing Configuration Guide*

 Copyright © 2018, Juniper Networks, Inc.

optics-options

Syntax

```
optics-options {
  alarm low-light-alarm {
    (link-down | syslog);
    }
  tca tca-identifier (enable-tca | no-enable-tca) (threshold number | threshold-24hrs number);
  tx-power dbm;
  warning low-light-warning {
    (link-down | syslog);
  }
  wavelength nm;
}
```

Hierarchy Level

[edit interfaces *interface-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
alarm option and **warning** options introduced in Junos OS Release 10.0.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement and **tx-power** option introduced in Junos OS Release 13.2 for PTX Series routers.
tca option introduced in Junos OS Release 14.2 for PTX Series routers.

Description

For 10-Gigabit Ethernet or 100-Gigabit Ethernet dense wavelength-division multiplexing (DWDM) interfaces only, configure full C-band International Telecommunication Union (ITU)-Grid tunable optics.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Ethernet DWDM Interface Wavelength Overview*
- *100-Gigabit Ethernet OTN Options Configuration Overview*

option-82

Syntax

```
option-82 <circuit-id> <remote-id>;
```

Hierarchy Level

[edit interfaces *interface-name* auto-configure vlan-ranges authentication username-include],
[edit interfaces *interface-name* auto-configure stacked-vlan-ranges authentication username-include]

Release Information

Statement introduced in Junos OS Release 10.0.
Options **circuit-id** and **remote-id** introduced in Junos OS Release 11.4.

Description

Specify that the option 82 information from the client PDU is concatenated with the username during the subscriber authentication process.

For autosense VLANs, you can additionally specify Option 82 suboption information that is concatenated with the username. You can specify either both or neither of the Agent Circuit ID (suboption 1) and Agent Remote ID (suboption 1). If you specify both, the Agent Circuit ID is supplied first, followed by a delimiter, and then the Agent Remote ID. If you specify that neither suboption is supplied, the raw payload of Option 82 from the PDU is concatenated to the username.

NOTE: **The option 82 value used in creating the username is based on the option 82 value that is encoded in the incoming DHCP discover packet. The use of suboptions is supported for DHCPv4 only.**

Options

none—Use the raw payload of Option 82 from the PDU.

circuit-id—(Optional) Use the Agent Circuit ID suboption (suboption 1).

remote-id—(Optional) Use the Agent Remote ID suboption (suboption 2).

Required Privilege Level

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

Related Documentation

- *Configuring VLAN Interface Username Information for AAA Authentication*
- *Using DHCP Option 82 Suboptions in Authentication Usernames for Autosense VLANs*

Copyright © 2018, Juniper Networks, Inc.

otn-options

Syntax

```
otn-options {
  bytes (otn-options) transmit-payload-type value;
  fec (efec | gfec | gfec-sdfec | none );
  (is-ma | no-is-ma);
  (laser-enable | no-laser-enable);
  (line-loopback | no-line-loopback);
  (local-loopback | no-local-loopback);
  (odu-ttim-action-enable | no-odu-ttim-action-enable);
  (otu-ttim-action-enable | no-otu-ttim-action-enable);
  odu-delay-management {
    (bypass | no-bypass);
    (monitor-end-point | no-monitor-end-point);
    number-of-frames value;
    (no-start-measurement | start-measurement;
  }
  odu-signal-degrade {
    ber-threshold-clear value;
    ber-threshold-signal-degrade value;
    interval value;
  }
  (prbs | no-prbs);
  preemptive-fast-reroute {
    (backward-frr-enable | no-backward-frr-enable);
    (signal-degrade-monitor-enable | no-signal-degrade-monitor-enable);
    odu-backward-frr-enable | no-odu-backward-frr-enable;
    odu-signal-degrade-monitor-enable | no-odu-signal-degrade-monitor-enable;
  }
  rate {
    (fixed-stuff-bytes | no-fixed-stuff-bytes);
    oc192;
    otu4;
    (pass-through | no-pass-through);
  }
  signal-degrade {
    ber-threshold-clear value;
    ber-threshold-signal-degrade value;
    interval value;
  }
  tca tca-identifier (enable-tca | no-enable-tca) (threshold number | threshold-24hrs number);
  transport-monitoring;
  trigger trigger-identifier;
  tti tti-identifier;
}
```

Hierarchy Level

[edit interfaces ge-*fpc/pic/port*]
[edit interfaces xe-fpc/pic/port]
[edit interfaces et-fpc/pic/port]

Release Information

Statement introduced in Junos OS Release 9.4.

bytes, **is-ma**, **local-loopback**, **no-is-ma**, **no-local-loopback**, **no-odu-ttim-action-enable**, **no-otu-ttim-action-enable**, **no-prbs**, **odu-delay-management**, **odu-ttim-action-enable**,

otu-ttim-action-enable, **prbs**, **preemptive-fast-reroute**, and **signal-degrade** statements introduced in Junos OS Release 13.2 for PTX Series routers.
oc192 statement introduced in Junos OS Release 13.3R3 for MX Series routers.
odu-signal-degrade, **odu-backward-frr-enable | no-odu-backward-frr-enable**, **odu-signal-degrade-monitor-enable | no-odu-signal-degrade-monitor-enable** statements introduced in Junos OS Release 14.1R2 and 14.2 for P2-100GE-OTN PIC in PTX5000 routers.
tca option introduced in Junos OS Release 14.2 for PTX Series routers.

Description Specify the Ethernet optical transport network (OTN) interface and options.

Options The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *10-Gigabit Ethernet OTN Options Configuration Overview*
- *100-Gigabit Ethernet OTN Options Configuration Overview*
- *Configuring 100-Gigabit DWDM OTN PICs*

output

Syntax

```
output {
  service-set service-set-name <service-filter filter-name>;
}
```

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Define one or more output service sets and filters to be applied to traffic.

Options The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

 Copyright © 2018, Juniper Networks, Inc.

output-list

Syntax output-list [*filter-names*];

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

Release Information Statement introduced in Junos OS Release 7.6.

Description Apply a group of filters to evaluate when packets are transmitted on an interface.

Options **[*filter-names*]**—Name of a filter to evaluate when packets are transmitted on the interface. Up to 16 filters can be included in a filter input list.

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

Related Documentation

- Applying a Filter to an Interface on page 235
- **input-list on page 675**
- *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*
- *Junos OS Services Interfaces Library for Routing Devices*
- *Junos OS Administration Library*

Copyright © 2018, Juniper Networks, Inc.

output-policer

Syntax output-policer *policer-name*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* layer2-policer],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* layer2-policer]

Release Information Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Apply a single-rate two-color policer to the Layer 2 output traffic at the logical interface. The **output-policer** and **output-three-color** statements are mutually exclusive.

Options *policer-name*—Name of the single-rate two-color policer that you define at the **[edit firewall]** hierarchy level.

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

Related Documentation

- *Two-Color and Three-Color Policers at Layer 2*
- *Applying Layer 2 Policers to Gigabit Ethernet Interfaces*
- *Configuring a Gigabit Ethernet Policer*
- input-policer on page 676
- input-three-color on page 678
- layer2-policer on page 718
- logical-interface-policer on page 747
- output-three-color on page 846

 Copyright © 2018, Juniper Networks, Inc.

output-priority-map

Syntax

```
output-priority-map {
  classifier {
    premium {
      forwarding-class class-name {
        loss-priority (high | low);
      }
    }
  }
}
```

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile
  ethernet-policer-profile]
[edit interfaces *interface-name* ether-options ethernet-switch-profile
  ethernet-policer-profile]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2 for the QFX Series.

Description

For Gigabit Ethernet IQ and 10-Gigabit Ethernet interfaces only, define the output policer priority map to be applied to outgoing frames on this interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Specifying an Output Priority Map*
- input-priority-map on page 677

output-three-color

Syntax output-three-color *policer-name*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* layer2-policer]
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* layer2-policer]

Release Information Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Apply a single-rate or two-rate three-color policer to the Layer 2 output traffic at the logical interface. The **output-three-color** and **output-policer** statements are mutually exclusive.

Options ***policer-name***—Name of the single-rate or two-rate three-color policer.

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

Related Documentation

- *Two-Color and Three-Color Policers at Layer 2*
- *Applying Layer 2 Policers to Gigabit Ethernet Interfaces*
- *Configuring a Gigabit Ethernet Policer*
- input-three-color on page 678
- input-policer on page 676
- layer2-policer on page 718
- logical-interface-policer on page 747
- output-policer on page 844

 Copyright © 2018, Juniper Networks, Inc.

output-vlan-map (Aggregated Ethernet)

Syntax

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

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 8.2.
Starting in Junos OS Release 17.3R1, input-vlan-map for outer vlan is supported for L2 circuit over aggregated Ethernet interfaces for QFX10000 Series switches.

Description

Define the rewrite profile to be applied to outgoing frames on this logical interface. On MX Series routers, this statement only applies to aggregated Ethernet inferfaces using Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2 and IQ2-E interfaces and 100-Gigabit Ethernet Type 5 PIC with CFP..

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Stacking and Rewriting Gigabit Ethernet VLAN Tags*
- input-vlan-map (Aggregated Ethernet) on page 679

output-vlan-map

Syntax

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

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.
pop-pop, **pop-swap**, **push-push**, **swap-push**, and **swap-swap** statements added in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

For EX Series switches, defines the rewrite operation to be applied to outgoing frames.

For MX Series routers and NFX Series devices' Gigabit Ethernet IQ and 10-Port 10-Gigabit Ethernet SFPP interfaces only, defines the rewrite operation to be applied to outgoing frames on this logical interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Stacking and Rewriting Gigabit Ethernet VLAN Tags*
- input-vlan-map on page 680

 Copyright © 2018, Juniper Networks, Inc.

overflow (Receive Bucket)

Syntax overflow (discard | tag);

Hierarchy Level [edit interfaces *interface-name* receive-bucket]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify how to handle packets that exceed the threshold for the receive leaky bucket.

Options **tag**—Tag, count, and process received packets that exceed the threshold.

discard—Discard received packets that exceed the threshold. No counting is done.

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

Related Documentation

- Configuring Receive and Transmit Leaky Bucket Properties to Reduce Network Congestion on page 156

overflow (Transmit Bucket)

Syntax overflow discard;

Hierarchy Level [edit interfaces *interface-name* transmit-bucket]

Release Information Statement introduced before Junos OS Release 7.4.

Description Discard packets that exceed the threshold for the transmit leaky bucket.

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

Related Documentation

- Configuring Receive and Transmit Leaky Bucket Properties to Reduce Network Congestion on page 156

Copyright © 2018, Juniper Networks, Inc.

override

Syntax override tag *vlan-tag* dynamic-profile *profile name*;

Hierarchy Level [edit interfaces *interface-name* auto-configure vlan-ranges],
[edit interfaces *interface-name* auto-configure stacked-vlan-ranges]

Release Information Statement introduced in Junos OS Release 11.2.

Description Override dynamic profile assignment to individual VLANs that are already part of a previously defined VLAN range and dynamic profile.

Options ***vlan-tag***—VLAN tag that you want to override.

profile-name—Name of the dynamic profile that you want to use when overriding the specified VLAN tag.

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

Related Documentation

- *Overriding the Dynamic Profile Used for an Individual VLAN*
- *Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs*
- *Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs*

 Copyright © 2018, Juniper Networks, Inc.

pado-advertise

Syntax pado-advertise;

Hierarchy Level [edit protocols pppoe]

Release Information Statement introduced in Junos OS Release 10.2.

Description Enable named services configured in PPPoE service name tables to be advertised in PPPoE Active Discovery Offer (PADO) control packets. By default, advertisement of named services in PADO packets is disabled.

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

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

Related Documentation

- *Configuring PPPoE Service Name Tables*
- *Enabling Advertisement of Named Services in PADO Control Packets*

paired-group

Syntax paired-group *group-name*;

Hierarchy Level [edit interfaces *interface-name* sonet-options aps]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure load sharing between two working protect circuit pairs.

Options *group-name*—Circuit's group name, as configured with the **protect-circuit** or **working-circuit** statement.

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

Related Documentation

- *Configuring APS Load Sharing*
- working-circuit on page 1106

Copyright © 2018, Juniper Networks, Inc.

pap

Syntax

```
pap {
  access-profile name;
  default-pap-password password;
  local-name name;
  local-password password;
  passive;
}
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 8.3.

Description

Configure the Password Authentication Protocol (PAP). Use PAP authentication as a means to provide a simple method for the peer to establish its identity using a two-way handshake. This is done only upon initial link establishment.

After the link is established, an ID and password pair is repeatedly sent by the peer to the authenticator until authentication is acknowledged or the connection is terminated.

BEST PRACTICE: On inline service (si) interfaces for L2TP, only the pap statement itself is typically used for subscriber management. We recommend that you leave the subordinate statements at their default values.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

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

Related Documentation

- Configuring the PPP Challenge Handshake Authentication Protocol on page 136
- Configuring the PPP Password Authentication Protocol On a Logical Interface on page 196
- Tracing Operations of the pppd Process on page 153
- traceoptions (PPP Process) on page 1033
- *Example: Configuring PAP for an L2TP Profile*
- *Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface*

 Copyright © 2018, Juniper Networks, Inc.

pap-password

Syntax pap-password *password*;

Hierarchy Level [edit access profile *profile-name* client *client-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure the Password Authentication Protocol (PAP) password.

NOTE: This statement is not supported for L2TP LNS on MX Series routers.

Options *password*—PAP password.

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

Related Documentation

- *Configuring the PAP Password for an L2TP Profile*

Copyright © 2018, Juniper Networks, Inc.

partition

Syntax

partition *partition-number* oc-slice *oc-slice-range* interface-type *type* timeslots *time-slot-range*;

Hierarchy Level

[edit interfaces *interface-name*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

The partition number is correlated with the channel number. Partition and channel numbering on IQ interfaces begins with **:1**, not **:0**.

Default

If you omit this statement, the channelized PIC or PIM is not partitioned, and no data channels are configured.

Options

partition-number—Sublevel interface partition index.

Range:

- 1 through 4 for an OC3 interface on a channelized OC12 IQ interface.
- 1 through 12 for a T3 interface on a channelized OC12 IQ interface.
- 1 through 4 for a T3 interface on a channelized T3 IQ interface.
- 1 through 28 for a T1 IQ interface on a channelized OC12 IQ or channelized T3 IQ interface.
- 1 through 10 for an E1 interface on a channelized E1 IQ interface.
- 1 through 30 on a channelized E1 interface.
- 1 through 23 on a channelized T1 interface.
- 1 through 24 for *N*xDS0 interfaces on either channelized OC12 IQ or channelized DS3 IQ interfaces.
- 0 through 31(with 0 reserved for framing) for *N*xDS0 interfaces on channelized E1 IQ interfaces.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Channelized E1 IQ and IQE Interfaces Overview*
- *Channelized OC12/STM4 IQ and IQE Interfaces Overview*
- *Configuring Channelized T3 IQ Interfaces*
- no-partition on page 821

 Copyright © 2018, Juniper Networks, Inc.

passive (CHAP)

Syntax passive;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Do not challenge the peer, but respond if challenged. If you omit this statement from the configuration, the interface always challenges its peer.

For ATM2 IQ interfaces only, you can configure CHAP on the logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:

- **atm-ppp-llc**—PPP over AAL5 LLC encapsulation.
- **atm-ppp-vc-mux**—PPP over AAL5 multiplex encapsulation.

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

Related Documentation
- *Configuring Passive Mode*

passive (PAP)

Syntax passive;

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

Release Information Statement introduced in Junos OS Release 8.3.

Description Initiate an authentication request when the PAP option is received from a peer. If you omit this statement from the configuration, the interface requires the peer to initiate an authentication request.

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

Related Documentation

- *Configuring Passive Mode*
- *Junos OS Administration Library*

 Copyright © 2018, Juniper Networks, Inc.

passive-monitor-mode

Syntax

passive-monitor-mode;

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Monitor packet flows from another router. If you include this statement in the configuration, the interface does not send keepalives or alarms, and does not participate actively on the network.

This statement is supported on ATM, Ethernet, and SONET/SDH interfaces. For more information, see *ATM Interfaces Feature Guide for Routing Devices*.

For ATM and Ethernet interfaces, you can include this statement on the physical interface only.

For SONET/SDH interfaces, you can include this statement on the logical interface only.

Required Privilege Level

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

Related Documentation

- *Enabling Passive Monitoring on ATM Interfaces*
- *Passive Monitoring on Ethernet Interfaces Overview*
- *Enabling Packet Flow Monitoring on SONET/SDH Interfaces*
- multiservice-options on page 795
- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2018, Juniper Networks, Inc.

password (Interfaces)

Syntax password *password-string*;

Hierarchy Level [edit interfaces *interface-name* auto-configure vlan-ranges authentication],
[edit interfaces *interface-name* auto-configure stacked-vlan-ranges authentication]

Release Information Statement introduced in Junos OS Release 10.0.

Description Configure the password that is sent to the external AAA authentication server for subscriber VLAN or stacked VLAN interface authentication.

Options ***password-string***—Authentication password.

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

Related Documentation

- *Configuring an Authentication Password for VLAN or Stacked VLAN Ranges*

path-database-size

Syntax path-database-size *path-database-size*;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management linktrace]

Release Information Statement introduced in Junos OS Release 8.5.

Description Number of linktrace reply entries to be stored per linktrace request.

Options **path-database-size**—Database size.

Range: 1 through 255

Default: 64

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

Related Documentation

- *Configuring Linktrace Protocol in CFM*

 Copyright © 2018, Juniper Networks, Inc.

path-trace

Syntax path-trace *trace-string*;

Hierarchy Level [edit interfaces *interface-name* sonet-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For SONET/SDH interfaces and 10-Gigabit Ethernet interfaces in WAN PHY mode, configure a path trace identifier, which is a text string that identifies the circuit.

On SONET/SDH OC48 interfaces that are configured for channelized (multiplexed) mode (by including the **no-concatenate** statement at the **[edit chassis fpc *slot-number* pic *pic-number*]** hierarchy level), the **bytes e1-quiet** and **bytes f1** options have no effect. The **bytes f2**, **bytes z3**, **bytes z4**, and **path-trace** options work correctly on channel 0 and work in the transmit direction only on channels 1, 2, and 3.

For DS3 channels on a channelized OC12 interface, you can configure a unique path trace for each of the 12 channels. Each path trace can be up to 16 bytes. For channels on a channelized OC12 IQ interface, each path trace can be up to 64 bytes.

Options ***trace-string***—Text string that identifies the circuit. If the string contains spaces, enclose it in quotation marks. A common convention is to use the circuit identifier as the path trace identifier. If you do not configure an identifier, the Junos OS uses the system and interface names to construct the default ***trace-string***. For all nonchannelized SONET/SDH interfaces, the default ***trace-string*** is ***system-name interface-name***. For channelized SONET/SDH interfaces and 10–Gigabit Ethernet WAN-PHY interfaces, the default ***trace-string*** is ***interface-name***.

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

Related Documentation

- *Configuring the SONET/SDH Path Trace Identifier for a Circuit*
- sonet-options on page 973

payload-scrambler

Syntax (payload-scrambler | no-payload-scrambler);

Hierarchy Level [edit interfaces *interface-name* e3-options],
[edit interfaces *interface-name* sonet-options],
[edit interfaces *interface-name* t3-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Enable or disable HDLC scrambling on an E3, a SONET/SDH, or a T3 interface. This type of scrambling provides better link stability. Both sides of a connection must either use or not use scrambling.

If you commit a T3 interface configuration that has HDLC payload scrambling enabled, the interface must also be configured to be compatible with the channel service unit (CSU) at the remote end of the line.

Disable payload scrambling on an E3 interface if Digital Link compatibility mode is used.

On a channelized OC12 interface, the **sonet payload-scrambler** statement is ignored. To configure scrambling on the DS3 channels on the interface, you can include the **t3-options payload-scrambler** statement in the configuration for each DS3 channel.

NOTE: The **payload-scrambler** statement at the **[edit interfaces** ***interface-name*** **e3-options]** hierarchy level is not valid for IQE PICs.

Default Payload scrambling is disabled on all E3 and T3 interfaces; it is enabled by default on E3/T3 over ATM interfaces and on SONET/SDH interfaces.

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

Related Documentation

- *Configuring E3 and T3 Parameters on ATM Interfaces*
- *Configuring E3 HDLC Payload Scrambling*
- *Configuring SONET/SDH HDLC Payload Scrambling for Link Stability*
- *Configuring T3 HDLC Payload Scrambling*
- *Examples: Configuring T3 Interfaces*
- compatibility-mode on page 491

 Copyright © 2018, Juniper Networks, Inc.

payload-size

Syntax payload-size *bytes* ;

Hierarchy Level [edit interfaces *interface-name* satop-options]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the **satop-options** payload size in integer number of bytes.

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

Related Documentation

- *ATM Support on Circuit Emulation PICs Overview*
- satop-options on page 948

pdu-interval

Syntax pdu-interval *interval*;

Hierarchy Level [edit protocols oam ethernet link-fault-management interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.2 for MX, M, T, ACX, Series routers, SRX Series firewalls, and EX Series Switches.
Statement introduced in Junos OS Release 9.4 for EX Series switches.

Description For Ethernet interfaces on EX Series switches and M320, M120, MX Series, and T Series routers, specify the periodic OAM PDU sending interval for fault detection. Used for IEEE 802.3ah Operation, Administration, and Management (OAM) support.

Options *interval*—Periodic OAM PDU sending interval.

Range: For MX, M, T, ACX, Series routers, SRX Series firewalls and EX Series switches – 100 through 1000 milliseconds

Default: For EX Series switches –1000 milliseconds

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

Related Documentation

- *Configuring the OAM PDU Interval*
- *Example: Configuring Ethernet OAM Link Fault Management*
- *Configuring Ethernet OAM Link Fault Management (CLI Procedure)*

 Copyright © 2018, Juniper Networks, Inc.

pdu-threshold

Syntax pdu-threshold *threshold-value*;

Hierarchy Level [edit protocols oam ethernet link-fault-management interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.2 for T, M, MX and ACX Series routers, SRX Series firewalls and EX Series switches.
Statement introduced in Junos OS Release 9.4 for EX Series switches and NFX Series devices.

Description Configure how many protocol data units (PDUs) are missed before declaring the peer lost in Ethernet OAM link fault management (LFM) for all interfaces or for specific interfaces.

For Ethernet interfaces on EX Series switches and M320, M120, MX Series, and T Series routers, specify the number of OAM PDUs to miss before an error is logged. Used for IEEE 802.3ah Operation, Administration, and Management (OAM) support.

Options ***threshold-value***—The number of PDUs missed before declaring the peer lost.

Range: 3 through 10 PDUs

Default: 3 PDUs

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

Related Documentation

- *Configuring the OAM PDU Threshold*
- *Configuring Ethernet OAM Link Fault Management (CLI Procedure)*

Copyright © 2018, Juniper Networks, Inc.

peer

Syntax

peer *a.b.c.d* {
  interface *interface-name*;
}

Hierarchy Level

[edit interfaces *interface-name* multi-chassis-protection]

Release Information

Statement introduced in Junos OS Release 11.1.

Description

For MX Series routers with multichassis aggregated Ethernet (MC-AE) interfaces, use the **multi-chassis-protection** statement under the physical interface level to reduce the configuration at the logical interface level. If the interchassis control protocol connection (ICCP) is UP and the interchassis data link (ICL) comes UP, the router configured as standby will bring up the MC-AE interfaces shared with the peer active-active node specified by the **peer** statement. You must also specify the peer's physical interface.

Options

a.b.c.d—Specify the IP address of the peer.

interface *interface-name*—Specify the peer's physical interface: interface *interface-name-fpc/pic/port*

Required Privilege Level

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

Related Documentation

- *Configuring Multichassis Link Aggregation on MX Series Routers*
- *Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation*
- *Configuring Aggregated Ethernet Link Protection*
- *Example: Configuring Aggregated Ethernet Link Protection*
- multi-chassis-protection on page 789

 Copyright © 2018, Juniper Networks, Inc.

peer-unit

Syntax peer-unit *unit-number*;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure a peer relationship between two logical systems.

Options ***unit-number***—Peering logical system unit number.

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

per-unit-scheduler

Syntax per-unit-scheduler;

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2 on 16x10GE MPC and MPC3E line cards.
Statement introduced in Junos OS Release 13.2 on PTX Series Packet Transport Routers.
Statement introduced in Junos OS Release 13.3 on MPC4E line cards.
Statement introduced in Junos OS Release 15.1 on MPC6E line cards.

Description For Channelized OC3 IQ, Channelized OC12 IQ, Channelized STM1 IQ, Channelized T3 IQ, Channelized E1 IQ, E3 IQ, link services IQ interfaces (lsq-), Gigabit Ethernet IQ, Gigabit Ethernet IQ2 and IQ2-E, and 10-, 40-, and 100-Gigabit Ethernet interfaces (including the 16x10GE MPC), enable the association of scheduler maps with logical interfaces.

CAUTION: Turning on per-unit scheduling causes the interface to reinitialize, which means all logical interfaces (units) on the interface are deleted and recreated.

NOTE: To enable per-unit scheduling on MX80 and MX104 routers, configure the **per-unit-scheduler** statement at each member physical interface level of a particular aggregated Ethernet interface as well as at that aggregated Ethernet interface level. On other routing platforms, it is enough if you include this statement at the aggregated Ethernet interface level.

NOTE: Per-unit scheduling is not supported on T1 interfaces configured on the Channelized OC12 IQ PIC.

NOTE: On Gigabit Ethernet IQ2 and IQ2-E PICs without the **per-unit-scheduler** statement, the entire PIC supports 4071 VLANs and the user can configure all the VLANs on the same port.

On Gigabit Ethernet IQ2 and IQ2-E PICs with the **per-unit-scheduler** statement, the entire PIC supports 1024 – 2 * number of ports (1024 minus two times the number of ports), because each port is allocated two default schedulers.

 Copyright © 2018, Juniper Networks, Inc.

When including the **per-unit-scheduler** statement, you must also include the **vlan-tagging** statement or the **flexible-vlan-tagging** statement (to apply scheduling to VLANs) or the **encapsulation frame-relay** statement (to apply scheduling to DLCIs) at the **[edit interfaces *interface-name*]** hierarchy level.

Required Privilege Level

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

Related Documentation

- *Applying Scheduler Maps and Shaping Rate to DLCIs and VLANs*
- vlan-tagging on page 1089
- flexible-vlan-tagging on page 618
- *Example: Applying Scheduling and Shaping to VLANs*
- *Configuring Virtual LAN Queuing and Shaping on PTX Series Routers*

Copyright © 2018, Juniper Networks, Inc.

performance-monitoring

Syntax

```
performance-monitoring {
  delegate-server-processing;
  hardware-assisted-timestamping;
  hardware-assisted-keepalives;
  sla-iterator-profiles {
    profile-name {
      avg-fd-twoway-threshold;
      avg-ifdv-twoway-threshold;
      avg-flr-forward-threshold;
      avg-flr-backward-threshold;
      disable;
      calculation-weight {
        delay delay-weight;
        delay-variation delay-variation-weight;
      }
      cycle-time milliseconds;
      iteration-period connections;
      measurement-type (loss | statistical-frame-loss | two-way-delay);
    }
  }
}
```

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management]

Release Information Statement introduced in Junos OS Release 9.5.

Description Specify performance monitoring support for Ethernet frame delay measurement.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- *Ethernet Frame Delay Measurements Overview*
- *Guidelines for Configuring Routers to Support an ETH-DM Session*
- *Enabling the Hardware-Assisted Timestamping Option*

 Copyright © 2018, Juniper Networks, Inc.

periodic

List of Syntax

Syntax (EX Series) on page 869
Syntax (QFX Series) on page 869

Syntax (EX Series)

periodic *interval*;

Syntax (QFX Series)

periodic (fast | slow);

Hierarchy Level (EX Series)

[edit interfaces ae*x* aggregated-ether-options lacp],
[edit interfaces interface-range *name* aggregated-ether-options lacp]

Hierarchy Level (QFX Series)

[edit interfaces ae*x* aggregated-ether-options lacp]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.

Description

For aggregated Ethernet interfaces only, configure the interval for periodic transmission of LACP packets.

Options

interval—Interval for periodic transmission of LACP packets.

- **fast**—Transmit packets every second.
- **slow**—Transmit packets every 30 seconds.

Default: fast

Required Privilege Level

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

Related Documentation

- *Configuring LACP for Aggregated Ethernet Interfaces*
- *Configuring Aggregated Ethernet LACP (CLI Procedure)*
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*
- *Configuring Aggregated Ethernet LACP (CLI Procedure)*
- *Understanding Aggregated Ethernet Interfaces and LACP for Switches*
- *Junos OS Network Interfaces Library for Routing Devices*

Copyright © 2018, Juniper Networks, Inc.

pfc

Syntax pfc;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description For interfaces with PPP encapsulation, configure the router to compress the protocol field to one byte.

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

Related Documentation
- Configuring the PPP Protocol Field Compression on page 151

pic-type

Syntax pic-type (atm1 | atm2);

Hierarchy Level [edit interfaces at-*fpc/pic/port* atm-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For ATM interfaces, configure the type of ATM PIC installed in your router.

Options **atm1**—ATM1 PIC.

atm2—ATM2 IQ PIC.

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

Related Documentation
- *Configuring the ATM PIC Type*

 Copyright © 2018, Juniper Networks, Inc.

plp-to-clp

Syntax
plp-to-clp;

Hierarchy Level
[edit interfaces at-*fpc/pic/port* atm-options],
[edit interfaces at-*fpc/pic/port* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces at--*fpc/pic/port* unit *logical-unit-number*]

Release Information
Statement introduced before Junos OS Release 7.4.

Description
For ATM2 IQ interfaces only, enable the PLP setting to be copied to the cell-loss priority (CLP) bit.

Default
If you omit this statement, the Junos OS does not copy the PLP setting to the CLP bit.

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

Related Documentation
- *Enabling the PLP Setting to Be Copied to the CLP Bit*
- *Copying the Packet Loss Priority to the CLP Bit on ATM Interfaces*

Copyright © 2018, Juniper Networks, Inc.

plp1

Syntax plp1 *cells*;

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

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.1 for QFX Series switches.

Description For ATM2 IQ interfaces only, define the EPD threshold on a VC. The EPD threshold is a limit on the number of transmit packets that can be queued. Packets that exceed the limit are discarded. This threshold applies to packets that have a PLP of 1.

Default EPD threshold is unregulated.

Options *cells*—Maximum number of cells.

Range: For 1-port and 2-port OC12 interfaces, 1 through 425,984 cellsFor 1-port OC48 interfaces, 1 through 425,984 cellsFor 2-port OC3, DS3, and E3 interfaces, 1 through 212,992 cellsFor 4-port DS3 and E3 interfaces, 1 through 106,496 cells

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

Related Documentation

- *Configuring Two EPD Thresholds per Queue*
- *Configuring an ATM Scheduler Map*
- linear-red-profile on page 725

 Copyright © 2018, Juniper Networks, Inc.

point-to-point

Syntax

```
point-to-point;
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

For all interfaces except aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet, configure the interface unit as a point-to-point connection. This is the default connection type.

Default

If you omit this statement, the interface unit is configured as a point-to-point connection.

Required Privilege Level

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

Related Documentation

- Configuring a Point-to-Point Connection on page 193
- multipoint on page 793

policer (CFM Firewall)

Syntax

```
policer cfm-policer {
  if-exceeding {
    bandwidth-limit 8k;
    burst-size-limit 2k;
  }
  then discard;
}
```

Hierarchy Level

[edit firewall]

Release Information

Statement introduced in Junos OS Release 10.0.

Description

Attach an explicit policer to CFM sessions.

Required Privilege Level

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

Related Documentation

- *Configuring Rate Limiting of Ethernet OAM Messages*

- policer (CFM Global) on page 874
- policer (CFM Session) on page 875

Copyright © 2018, Juniper Networks, Inc.

policer (CFM Global)

Syntax

```
policer {
  all cfm-policer-name;
  continuity-check cfm-policer-name;
  other cfm-policer-name;
  }
}
```

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management]

Release Information Statement introduced in Junos OS Release 10.0.

Description Specify a policer at the global level to police the CFM traffic belonging to all sessions.

Options

continuity-check *cfm-policer-name*—Police all continuity check packets with the policer specified.

other *cfm-policer-name*—Police all non-continuity check packets with the policer specified.

all *cfm-policer-name*—Police all CFM packets with policer specified. If the **all** option is used, then you cannot specify above two options.

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

Related Documentation

- *Configuring Rate Limiting of Ethernet OAM Messages*
- policer (CFM Session) on page 875

 Copyright © 2018, Juniper Networks, Inc.

policer (CFM Session)

Syntax

policer {
  all *cfm-policer-name*;
  continuity-check *cfm-policer-name*;
  other *cfm-policer-name*;
  }
}

Hierarchy Level

[edit protocols oam ethernet connectivity-fault-management maintenance-domain *name* level *number* maintenance-association *name*]

Release Information

Statement introduced in Junos OS Release 10.0.

Description

Specify a separate policer to rate-limit packets specific to that session.

Options

- **continuity-check** ***cfm-policer-name***—Police continuity check packets belonging to this session.

  other ***cfm-policer-name***—Police all non-continuity check packets belonging to this session.

  all ***cfm-policer-name***—Police all CFM packets belonging to this session. If the **all** option is used, then you cannot specify the above two options.

Required Privilege Level

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

Related Documentation

- *Configuring Rate Limiting of Ethernet OAM Messages*
- policer (CFM Global) on page 874

policer (CoS)

Syntax

```
policer cos-policer-name {
  aggregate {
    bandwidth-limit bps;
    burst-size-limit bytes;
  }
  premium {
    bandwidth-limit bps;
    burst-size-limit bytes;
  }
}
```

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

For Gigabit Ethernet IQ , Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), and 100-Gigabit Ethernet Type 5 PIC with CFP, define a CoS policer template to specify the premium bandwidth and burst-size limits, and the aggregate bandwidth and burst-size limits. The premium policer is not supported on MX Series routers or for Gigabit Ethernet interfaces with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router).

Options

cos-policer-name—Name of one policer to specify the premium bandwidth and burst-size limits, and the aggregate bandwidth and burst-size limits.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Configuring Gigabit Ethernet Policers*

 Copyright © 2018, Juniper Networks, Inc.

policer (Interface)

Syntax

```
policer {
  arp policer-template-name;
  input policer-template-name;
  output policer-template-name;
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Apply a policer to an interface.

Options

arp ***policer-template-name***—For **inet** family only, name of one policer to evaluate when ARP packets are received on the interface.

input ***policer-template-name***—Name of one policer to evaluate when packets are received on the interface.

output ***policer-template-name***—Name of one policer to evaluate when packets are transmitted 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

- Applying Policers on page 226
- *Configuring Firewall Filters and Policers for VPLS*
- *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*
- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2018, Juniper Networks, Inc.

policer (MAC)

Syntax

```
policer {
  input cos-policer-name;
  output cos-policer-name;
}
```

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* accept-source-mac mac-address *mac-address*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* accept-source-mac mac-address *mac-address*]

Release Information

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

Description

For Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), and 100-Gigabit Ethernet Type 5 PIC with CFP, configure MAC policing.

NOTE:
On MX Series routers with Gigabit Ethernet or Fast Ethernet PICs, the following considerations apply:

- Interface counters do not count the 7-byte preamble and 1-byte frame delimiter in Ethernet frames.
- In MAC statistics, the frame size includes MAC header and CRC before any VLAN rewrite/imposition rules are applied.
- In traffic statistics, the frame size encompasses the L2 header without CRC after any VLAN rewrite/imposition rule.

Options

input *cos-policer-name*—Name of one policer to specify the premium bandwidth and aggregate bandwidth.

output *cos-policer-name*—Name of one policer to specify the premium bandwidth and aggregate bandwidth.

Required Privilege Level

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

Related Documentation

- *Configuring MAC Address Filtering*

Copyright © 2018, Juniper Networks, Inc.

policy-statement

Syntax

```
policy-statement policy-name {
  term term-name {
    from {
      as-path-unique-count count (equal | orhigher | orlower);
      family family-name;
      match-conditions;
      policy subroutine-policy-name;
      prefix-list prefix-list-name;
      prefix-list-filter prefix-list-name match-type <actions>;
      protocol protocol-name;
      route-filter destination-prefix match-type <actions>;
      source-address-filter source-prefix match-type <actions>;
      tag value;
      traffic-engineering;
    }
    to {
      match-conditions;
      policy subroutine-policy-name;
    }
    then actions;
  }
  then {
    aggregate-bandwidth;
    dynamic-tunnel-attributes dynamic-tunnel-attributes;
    limit-bandwidth limit-bandwidth;
    multipath-resolve multipath-resolve;
    no-entropy-label-capability;
    prefix-segment {
      index index;
      node-segment;
    }
    priority (high | medium | low);
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* policy-options],
[edit logical-systems *logical-system-name* policy-options],
[edit policy-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for configuration in the dynamic database introduced in Junos OS Release 9.5.
Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.
inet-mdt option introduced in Junos OS Release 10.0R2.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
route-target option introduced in Junos OS Release 12.2.
Statement introduced in Junos OS 14.1X53-D20 for the OCX Series.
protocol and **traffic-engineering** options introduced in Junos OS Release 14.2.
no-entropy-label-capability option introduced in Junos OS Release 15.1.

Copyright © 2018, Juniper Networks, Inc.

priority and **tag** ***value*** options introduced in Junos OS Release 17.1.
as-path-unique-count option introduced in Junos OS Release 17.2R1.
prefix-segment option introduced in Junos OS Release 17.2R1 for MX Series routers, PTX Series routers, QFX5100 switches, and QFX10000 switches.
multipath-resolve and **dynamic-tunnel-attributes** options introduced in Junos OS Release 17.3R1.
aggregate-bandwidth and **limit-bandwidth** ***limit-bandwidth*** options introduced in Junos OS Release 17.4R1 for MX Series, PTX Series, and QFX Series.

Description Define a routing policy, including subroutine policies.

A *term* is a named structure in which match conditions and actions are defined. Routing policies are made up of one or more terms. Each routing policy term is identified by a term name. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks.

Each term contains a set of match conditions and a set of actions:

- Match conditions are criteria that a route must match before the actions can be applied. If a route matches all criteria, one or more actions are applied to the route.
- Actions specify whether to accept or reject the route, control how a series of policies are evaluated, and manipulate the characteristics associated with a route.

Generally, a router compares a route against the match conditions of each term in a routing policy, starting with the first and moving through the terms in the order in which they are defined, until a match is made and an explicitly configured or default action of **accept** or **reject** is taken. If none of the terms in the policy match the route, the router compares the route against the next policy, and so on, until either an action is taken or the default policy is evaluated.

If none of the match conditions of each term evaluates to true, the final action is executed. The final action is defined in an unnamed term. Additionally, you can define a default action (either **accept** or **reject**) that overrides any action intrinsic to the protocol.

The order of match conditions in a term is not relevant, because a route must match all match conditions in a term for an action to be taken.

To list the routing policies under the **[edit policy-options]** hierarchy level by **policy-statement** ***policy-name*** in alphabetical order, enter the **show policy-options** configuration command.

The statements are explained separately.

 Copyright © 2018, Juniper Networks, Inc.

Options *actions*—(Optional) One or more actions to take if the conditions match. The actions are described in *Configuring Flow Control Actions*.

family ***family-name***—(Optional) Specify an address family protocol. Specify **inet** for IPv4. Specify **inet6** for 128-bit IPv6, and to enable interpretation of IPv6 router filter addresses. For IS-IS traffic, specify **iso**. For IPv4 multicast VPN traffic, specify **inet-mvpn**. For IPv6 multicast VPN traffic, specify **inet6-mvpn**. For multicast-distribution-tree (MDT) IPv4 traffic, specify **inet-mdt**. For BGP route target VPN traffic, specify **route-target**. For traffic engineering, specify **traffic-engineering**.

NOTE: When **family** is not specified, the routing device or routing instance uses the address family or families carried by BGP. If multiprotocol BGP (MP-BGP) is enabled, the policy defaults to the protocol family or families carried in the network layer reachability information (NLRI) as configured in the *family* statement for BGP. If MP-BGP is not enabled, the policy uses the default BGP address family unicast IPv4.

from—(Optional) Match a route based on its source address.

as-path-unique-count ***count*** **(equal | orhigher | orlower)**—(Optional) Specify a number from 0 through 1024 to filter routes based on the number of unique autonomous systems (ASs) in the AS path. Specify the match condition for the unique AS path count.

aggregate-bandwidth—(Optional) Enable BGP to advertise aggregate outbound link bandwidth for load balancing.

dynamic-tunnel-attributes ***dynamic-tunnel-attributes***—(Optional) Choose a set of defined dynamic tunnel attributes for forwarding traffic over V4oV6 tunnels.

match-conditions—(Optional in **from** statement; required in **to** statement) One or more conditions to use to make a match. The qualifiers are described in *Routing Policy Match Conditions*.

multipath-resolve ***multipath-resolve***–(Optional) Enable the use of all paths for resolution over the specified prefix.

limit-bandwidth ***limit-bandwidth***—(Optional) Specify the limit for advertised aggregate outbound link bandwidth for load balancing.

Range: 0 through 4,294,967,295 bytes

no-entropy-label-capability—(Optional) Disable the entropy label capability advertisement at egress or transit routes specified in the policy.

priority (high | medium | low)—(Optional) Configure the priority for an IS-IS route to change the default order in which the routes are installed in the routing table, in the event of a network topology change.

Copyright © 2018, Juniper Networks, Inc.

policy ***subroutine-policy-name***—Use another policy as a match condition within this policy. The name identifying the subroutine policy can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (“ ”). Policy names cannot take the form **__.*-internal__**, as this form is reserved. For information about how to configure subroutines, see *Understanding Policy Subroutines in Routing Policy Match Conditions*.

policy-name—Name that identifies the policy. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (“ ”).

prefix-list ***prefix-list-name***—Name of a list of IPv4 or IPv6 prefixes.

prefix-list-filter ***prefix-list-name***—Name of a prefix list to evaluate using qualifiers; ***match-type*** is the type of match, and ***actions*** is the action to take if the prefixes match.

protocol ***protocol-name***—Name of the protocol used to control traffic engineering database import at the originating point.

route-filter ***destination-prefix match-type <actions>***—(Optional) List of routes on which to perform an immediate match; ***destination-prefix*** is the IPv4 or IPv6 route prefix to match, ***match-type*** is the type of match (see *Configuring Route Lists*), and ***actions*** is the action to take if the ***destination-prefix*** matches.

source-address-filter ***source-prefix match-type <actions>***—(Optional) Unicast source addresses in multiprotocol BGP (MBGP) and Multicast Source Discovery Protocol (MSDP) environments on which to perform an immediate match. ***source-prefix*** is the IPv4 or IPv6 route prefix to match, ***match-type*** is the type of match (see *Configuring Route Lists*), and ***actions*** is the action to take if the ***source-prefix*** matches.

tag ***value***—(Optional) A numeric value that identifies a route. You can tag certain routes to prioritize them over other routes. In the event of a network topology change, Junos OS updates these routes in the routing table before updating other routes with lower priority. You can also tag some routes to identify and reject them based on your requirement.

term ***term-name***—Name that identifies the term. The term name must be unique in the policy. It can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the entire name in quotation marks (“ ”). A policy statement can include multiple terms. We recommend that you name all terms. However, you do have the option to include an unnamed term which must be the final term in the policy. To configure an unnamed term, omit the **term** statement when defining match conditions and actions.

to—(Optional) Match a route based on its destination address or the protocols into which the route is being advertised.

then—(Optional) Actions to take on matching routes. The actions are described in *Configuring Flow Control Actions* and *Configuring Actions That Manipulate Route Characteristics*.

 Copyright © 2018, Juniper Networks, Inc.

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

Related Documentation

- *dynamic-db*
- *Understanding Source Packet Routing in Networking (SPRING)*

pool

Syntax pool *pool-name* <priority *priority*>;

Hierarchy Level [edit interfaces br-*pim*/0/*port* dialer-options],
[edit interfaces umd0 dialer-options],
[edit interfaces dl*n* unit *logical-unit-number* dialer-options],
[edit logical-systems *logical-system-name* interfaces dl*n* unit *logical-unit-number* dialer-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description On J Series Services Routers, for logical and physical ISDN interfaces, specify the dial pool. The dial pool allows logical (dialer) and physical (**br-*pim*/0/*port***) interfaces to be bound together dynamically on a per-call basis. On a dialer interface, **pool** directs the dialer interface which dial pool to use. On **br-*pim*/0/*port*** interface, **pool** defines the pool to which the interface belongs.

Options ***pool-name***—Pool identifier.

priority *priority*—(Physical **br-*pim*/0/*port*** interfaces only) Specify a priority value of 0 (lowest) to 255 (highest) for the interface within the pool.

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

Related Documentation

- *Junos OS Interfaces and Routing Configuration Guide*

pop

Syntax

pop;

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.

Description

NOTE: On EX4300 switches, pop is not supported at the [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map] hierarchy level.

For Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2, and IQ2-E interfaces; 10-Gigabit Ethernet LAN/WAN PIC; aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces; 100-Gigabit Ethernet Type 5 PIC with CFP; and Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces, specify the VLAN rewrite operation to remove a VLAN tag from the top of the VLAN tag stack. The outer VLAN tag of the frame is removed.

Required Privilege Level

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

Related Documentation

- *Removing a VLAN Tag*
- *Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support (CLI Procedure)*

Copyright © 2018, Juniper Networks, Inc.

pop-all-labels

Syntax

```
pop-all-labels {
  required-depth number;
}
```

Hierarchy Level

[edit interfaces *interface-name* atm-options mpls],
[edit interfaces *interface-name* sonet-options mpls],
[edit interfaces *interface-name* fastether-options mpls],
[edit interfaces *interface-name* gigether-options mpls]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description

For passive monitoring on ATM, SONET/SDH, Fast Ethernet, and Gigabit Ethernet interfaces only, removes up to two MPLS labels from incoming IP packets. For passive monitoring on T Series devices, removes up to five MPLS labels from incoming IP packets.

This statement has no effect on IP packets with more than two MPLS labels, or IP packets with more than five MPLS labels on T Series devices. Packets with MPLS labels cannot be processed by the Monitoring Services PIC; if packets with MPLS labels are forwarded to the Monitoring Services PIC, they are discarded.

The remaining statement is explained separately. See CLI Explorer.

Default

If you omit this statement, the MPLS labels are not removed, and the packet is not processed by the Monitoring Services PIC.

Required Privilege Level

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

Related Documentation

- *Removing MPLS Labels from Incoming Packets*
- *Enabling Packet Flow Monitoring on SONET/SDH Interfaces*
- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2018, Juniper Networks, Inc.

pop-pop

Syntax pop-pop;

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

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description For Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP, and for 10-Gigabit Ethernet SFP interfaces on EX Series switches, specify the VLAN rewrite operation to remove both the outer and inner VLAN tags of the frame.

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

Related Documentation

- *Removing the Outer and Inner VLAN Tags*

 Copyright © 2018, Juniper Networks, Inc.

pop-swap

Syntax pop-swap;

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

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify the VLAN rewrite operation to remove the outer VLAN tag of the frame, and replace the inner VLAN tag of the frame with a user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame.

You can use this statement on Gigabit Ethernet IQ, IQ2, IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP.

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

Related Documentation

- *Removing the Outer VLAN Tag and Rewriting the Inner VLAN Tag*

port

Syntax

```
port {
  minimum port-number;
  maximum port-number;
}
```

Hierarchy Level [edit interfaces vsp-*fpc/pic/port* unit *logical-unit-number* compression rtp]

Release Information Statement introduced before Junos OS Release 7.4.

Description For voice services interfaces only, assign User Datagram Protocol (UDP) destination port numbers reserved for Real-Time Transport Protocol (RTP) traffic.

Options **minimum *port-number***—Specify minimum port number.

Range: 0 through 65,535

maximum *port-number*—Specify maximum port number.

Range: 0 through 65,535

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

 Copyright © 2018, Juniper Networks, Inc.

port-priority

Syntax port-priority *priority*;

Hierarchy Level [edit interfaces *interface-name* gigether-options 802.3ad lacp]

Release Information Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 11.4 for EX Series switches.
Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.

Description Define LACP port priority at the interface level.

Options *priority*—Priority for being elected to be the active port and both collect and distribute traffic. A smaller value indicates a higher priority for being elected.

Range: 0 through 65535

Default: 127

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

Related Documentation

- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*
- *Configuring Aggregated Ethernet LACP (CLI Procedure)*

Copyright © 2018, Juniper Networks, Inc.

port-status-tlv

Syntax

port-status-tlv blocked;

Hierarchy Level

[edit protocols oam ethernet connectivity-fault-management action-profile *tlv-action* event]
[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* continuity-check]

Release Information

Statement introduced in Junos OS Release 9.6.

Description

Define an **action-profile** consisting of various events and the action. Based on values of **port-status-tlv** in the received CCM packets, specific action such as *interface-down* can be taken using action-profile options.

Options

blocked—When the incoming CCM packet contains port status TLV with value blocked, the action will be triggered for this action-profile.

Required Privilege Level

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

Related Documentation

- *Configuring a CFM Action Profile to Specify CFM Actions for CFM Events*
- *Configuring Remote MEP Action Profile Support*

post-service-filter

Syntax

post-service-filter *filter-name*;

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* family inet service input],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet service input]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Define the filter to be applied to traffic after service processing. The filter is applied only if a service set is configured and selected.

Options

filter-name—Identifier for postservice filter.

Required Privilege Level

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

 Copyright © 2018, Juniper Networks, Inc.

ppp-options

Syntax

```
ppp-options {
  authentication [ authentication-protocols ];
    mru size;
  mtu (size | use-lower-layer);
  chap {
    access-profile name;
    challenge-length minimum minimum-length maximum maximum-length;
    default-chap-secret name;
    local-name name;
    passive;
  }
  compression {
    acfc;
    pfc;
  }
  dynamic-profile profile-name;
  initiate-ncp (ip | ipv6 | dual-stack-passive)
  ipcp-suggest-dns-option;
  lcp-max-conf-req number
  lcp-restart-timer milliseconds;
  loopback-clear-timer seconds;
  ncp-max-conf-req number
  ncp-restart-timer milliseconds;
  on-demand-ip-address
  pap {
    access-profile name;
    default-pap-password password;
    local-name name;
    local-password password;
    passive;
  }
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

On interfaces with PPP encapsulation, configure PPP-specific interface properties.

For ATM2 IQ interfaces only, you can configure CHAP on the logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:

- **atm-ppp-llc**—PPP over AAL5 LLC encapsulation.
- **atm-ppp-vc-mux**—PPP over AAL5 multiplex encapsulation.

Copyright © 2018, Juniper Networks, Inc.

BEST PRACTICE: On inline service (si) interfaces for L2TP, only the **chap** and **pap** statements are typically used for subscriber management. We recommend that you leave the other statements subordinate to **ppp-options**—including those subordinate to **chap** and **pap**—at their default values.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

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

Related Documentation

- Configuring the PPP Challenge Handshake Authentication Protocol on page 136
- *Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface*

 Copyright © 2018, Juniper Networks, Inc.

pppoe-options

Syntax

```
pppoe-options {
  access-concentrator name;
  auto-reconnect seconds;
  (client | server);
  service-name name;
  underlying-interface interface-name;
  ppp-max-payload ppp-max-payload
}
```

Hierarchy Level

[edit interfaces pp0 unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number*]
[set interface *ppp interface*unit *logical-unit-number* ppp-max-payload *ppp-max-payload*],

Release Information

Statement introduced before Junos OS Release 7.4.
client Statement introduced in Junos OS Release 8.5.
server Statement introduced in Junos OS Release 8.5.
client Statement introduced in Junos OS Release 15.1X49-D100.

Description

Configure PPP over Ethernet-specific interface properties.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

The PPP-Max-Payload option allows you to override the default behavior of the PPPoE client by providing a maximum size that the PPP payload can support in both sending and receiving directions. The PPPoE server might allow the negotiation of an MRU larger than 1492 octets and the ability to use an MTU larger than 1500 octets.

Required Privilege Level

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

Related Documentation

- *Configuring a PPPoE Interface*

pppoe-underlying-options (Static and Dynamic Subscribers)

Syntax

pppoe-underlying-options {
  access-concentrator *name*;
  dynamic-profile *profile-name*;
  direct-connect
  duplicate-protection;
  max-sessions *number*;
  max-sessions-vsa-ignore;
  service-name-table *table-name*;
  short-cycle-protection <lockout-time-min *minimum-seconds*> <lockout-time-max
    maximum-seconds> <filter [aci]>;
}

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 10.0.

Description

Configure PPPoE-specific interface properties for the underlying interface on which the router creates a static or dynamic PPPoE logical interface. The underlying interface must be configured with PPPoE (**ppp-over-ether**) encapsulation.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

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

Related Documentation

- *Configuring PPPoE* (for static interfaces)
- *Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces*
- *Assigning a Service Name Table to a PPPoE Underlying Interface*

Copyright © 2018, Juniper Networks, Inc.

preferred

Syntax

preferred;

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description

Configure this address to be the preferred address on the interface. If you configure more than one address on the same subnet, the preferred source address is chosen by default as the source address when you initiate frame transfers to destinations on the subnet.

NOTE: The edit logical-systems hierarchy is not available on QFabric systems.

Default

The lowest-numbered address on the subnet is the preferred address.

Required Privilege Level

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

Related Documentation

- Configuring the Interface Address on page 207

Copyright © 2018, Juniper Networks, Inc.

preferred-source-address

Syntax preferred-source-address *address*;

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

Release Information Statement introduced in Junos OS Release 9.0.

Description For unnumbered Ethernet interfaces configured with a loopback interface as the donor interface, specify one of the loopback interface's secondary addresses as the preferred source address for the unnumbered Ethernet interface. Configuring the preferred source address enables you to use an IP address other than the primary IP address on some of the unnumbered Ethernet interfaces in your network.

Configuration of a preferred source address for unnumbered Ethernet interfaces is supported for the IPv4 and IPv6 address families.

Options ***address***—Secondary IP address of the donor loopback interface.

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

Related Documentation

- Configuring a Preferred Source Address for Unnumbered Ethernet or Demux Interfaces on page 218
- **address on page 424**
- *Junos OS Administration Library*

 Copyright © 2018, Juniper Networks, Inc.

preferred-source-address

Syntax

preferred-source-address *address*;

Hierarchy Level

[edit dynamic-profiles interfaces *interface-name* unit *logical-unit-number* family *family* unnumbered-address *interface-name*],
[edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family *family*],

Release Information

Statement introduced in Junos OS Release 9.2.
Support for the **$junos-preferred-source-address** and **$junos-preferred-source-ipv6-address** predefined variables introduced in Junos OS Release 9.6.

Description

For unnumbered Ethernet interfaces configured with a loopback interface as the donor interface, specify one of the loopback interface's secondary addresses as the preferred source address for the unnumbered Ethernet interface. Configuring the preferred source address enables you to use an IP address other than the primary IP address on some of the unnumbered Ethernet interfaces in your network. To configure the preferred source address dynamically, instead of using this statement, you must include the **$junos-preferred-source-address** predefined variable for IPv4 (family inet) addresses or the **$junos-preferred-source-ipv6-address** predefined variable for IPv6 (family inet6) addresses.

Configuration of a preferred source address for unnumbered Ethernet interfaces is supported for IPv4 and IPv6 address families.

NOTE: When you specify a static logical interface for the unnumbered interface in a dynamic profile that includes the **$junos-routing-instance** predefined variable, you must not configure a preferred source address, whether with the **$junos-preferred-source-address** predefined variable, the **$junos-preferred-source-ipv6-address** predefined variable, or the **preferred-source-address** statement. Configuring the preferred source address in this circumstance causes a commit failure.

Options

address—Secondary IP address of the donor loopback interface. Alternatively, use the **$junos-preferred-source-address** or the **$junos-preferred-source-ipv6-address** predefined variable to dynamically apply a preferred source address to the unnumbered Ethernet interface.

Required Privilege Level

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

Related Documentation

- Configuring an Unnumbered Interface on page 218
- *Junos OS Network Interfaces Library for Routing Devices*

Copyright © 2018, Juniper Networks, Inc.

- *Junos OS Administration Library*

premium (Hierarchical Policer)

Syntax

```
premium {
  if-exceeding {
    bandwidth-limit bandwidth;
    burst-size-limit burst;
  }
  then {
    discard;
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* firewall hierarchical-policer],
[edit firewall hierarchical-policer]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles ... hierarchical-policer *name*]** hierarchy level introduced in Junos OS Release 11.4.

Description

On M40e, M120, and M320 edge routers with FPC input as FFPC and FPC output as SFPC, and on MX Series, T320, T640, and T1600 edge routers with Enhanced Intelligent Queuing (IQE) PICs, T4000 routers with Type 5 FPC and Enhanced Scaling Type 4 FPC, specify a premium level for a hierarchical policer.

Options

Options are described separately.

Required Privilege Level

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

Related Documentation

- Applying Policers on page 226
- *Guidelines for Applying Traffic Policers*
- *Hierarchical Policer Configuration Overview*
- *Hierarchical Policers*
- *aggregate (Hierarchical Policer)*
- bandwidth-limit (Hierarchical Policer) on page 456
- burst-size-limit (Hierarchical Policer) on page 469
- *hierarchical-policer*
- if-exceeding (Hierarchical Policer) on page 653

 Copyright © 2018, Juniper Networks, Inc.

premium (Output Priority Map)

Syntax

```
premium {
  forwarding-class class-name {
    loss-priority (high | low);
  }
}
```

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile output-priority-map classifier]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

For Gigabit Ethernet IQ interfaces only, define the classifier for egress premium traffic.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Specifying an Output Priority Map*
- input-priority-map on page 677

premium (Policer)

Syntax

```
premium {
  bandwidth-limit bps;
  burst-size-limit bytes;
}
```

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile policer *cos-policer-name*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Define a policer to apply to nonpremium traffic.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Configuring Gigabit Ethernet Policers*
- aggregate (Gigabit Ethernet CoS Policer) on page 429
- ieee802.1p on page 652

Copyright © 2018, Juniper Networks, Inc.

preserve-interface

Syntax preserve-interface;

Hierarchy Level [edit interfaces *interface-name* sonet-options aps]

Release Information Statement introduced in Junos OS Release 7.6.

Description Provide link PIC replication, providing MLPPP link redundancy at the port level. This feature is supported with SONET APS and the following link PICs:

- Channelized OC3 IQ PIC
- Channelized OC12 IQ PIC
- Channelized STM1 IQ PIC

Link PIC replication provides the ability to add two sets of links, one from the active SONET PIC and the other from the standby SONET PIC, to the same bundle. If the active SONET PIC fails, links from the standby PIC are used without triggering link renegotiation. All the negotiated state is replicated from the active links to the standby links to prevent link renegotiation.

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

Related Documentation

- *Configuring Link PIC Redundancy*
- *Junos OS Services Interfaces Library for Routing Devices*

 Copyright © 2018, Juniper Networks, Inc.

primary (Address on Interface)

Syntax primary;

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

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure this address to be the primary address of the protocol on the interface. If the logical unit has more than one address, the primary address is used by default as the source address when packet transfer originates from the interface and the destination address does not indicate the subnet.

NOTE: The edit logical-systems hierarchy is not available on QFabric systems.

Default For unicast traffic, the primary address is the lowest non-127 (in other words, non-loopback) preferred address on the unit.

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

Related Documentation

- Configuring the Interface Address on page 207

Copyright © 2018, Juniper Networks, Inc.

primary (Interface for Router)

Syntax primary;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the primary interface for a device. By default, the multicast-capable interface with the lowest-index address is chosen as the primary interface. If there is no such interface, the point-to-point interface with the lowest-index address is chosen. Otherwise, any interface with an address can be picked. In practice, this means that, on the device, the **fxp0** or **em0** interface is picked by default. To configure a different interface to be the primary interface, you include this statement.

The *primary interface* for the router has the following characteristics:

- It is the interface through which the packets go out when you type a command such as ping 255.255.255.255—that is, a command that does not include an interface name (there is no interface ***type*-0/0/0.0** qualifier) and where the destination address does not imply any particular outgoing interface.
- It is the interface on which multicast applications running locally on the router, such as Session Announcement Protocol (SAP), perform group joins by default.
- It is the interface from which the default local address is derived for packets sourced out of an unnumbered interface if there are no non-127 addresses configured on the loopback interface, lo0.

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

Related Documentation

- Configuring Default, Primary, and Preferred Addresses and Interfaces on page 209

 Copyright © 2018, Juniper Networks, Inc.

primary (AS PIC or Multiservices PIC Interfaces)

Syntax primary *interface-name*;

Hierarchy Level [edit interfaces (rsp0 | rsp1) redundancy-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the primary AS PIC or MultiServices PIC interface.

Options ***interface-name***—The identifier for the AS PIC interface or MultiServices PIC interface, which must be of the form **sp-*fpc/pic/port***.

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

Related Documentation
- *Junos OS Services Interfaces Library for Routing Devices*

priority (OAM Connectivity-Fault Management)

Syntax priority *number*;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id*]

For EX Series Switches:

[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id*]

Release Information Statement introduced in Junos OS Release 8.4.

Description IEEE 802.1p priority bits used by the continuity check messages.

Options ***number***—Configure the IEEE 802.1p priority bits to be used in the VLAN header of the CFM packets.

Range: 0 through 7

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

Related Documentation
- *Configuring a MEP to Generate and Respond to CFM Protocol Messages*

Copyright © 2018, Juniper Networks, Inc.

priority (Schedulers)

Syntax priority (high | low);

Hierarchy Level [edit interfaces at-*fpc/pic/port* atm-options scheduler-maps *map-name* forwarding-class *class-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For ATM2 IQ interfaces only, assign queuing priority to a forwarding class.

Options **low**—Forwarding class has low priority.

high—Forwarding class has high priority.

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

Related Documentation

- *ATM2 IQ VC Tunnel CoS Components Overview*

 Copyright © 2018, Juniper Networks, Inc.

promiscuous-mode

Syntax

```
promiscuous-mode {
  vpi vpi-identifier;
}
```

Hierarchy Level

[edit interfaces *interface-name* atm-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers.

Description

For ATM interfaces with **atm-ccc-cell-relay** encapsulation, map all incoming cells from either an interface port or a VP to a single label-switched path (LSP) without restricting the VCI number. Promiscuous mode allows you to map traffic from all 65,535 VCIs to a single LSP, or from all 256 VPIs to a single LSP.

NOTE: In ACX Series routers, the statement supports only Inverse Multiplexing for ATM (IMA).

Options

vpi-identifier—Open this VPI in promiscuous mode.

Range: 0 through 255

Required Privilege Level

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

Related Documentation

- *Configuring ATM Cell-Relay Promiscuous Mode*
- vpi (ATM CCC Cell-Relay Promiscuous Mode) on page 1095

protect-circuit

Syntax protect-circuit *group-name*;

Hierarchy Level [edit interfaces *interface-name* sonet-options aps]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the protect router in an APS circuit pair. When the working interface fails, APS brings up the protection circuit and the traffic is moved to the protection circuit.

Options ***group-name***—Circuit's group name.

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

Related Documentation

- *Configuring Basic Automatic Protect Switching*
- working-circuit on page 1106

 Copyright © 2018, Juniper Networks, Inc.

protection-group

Syntax

```
protection-group {
  ethernet-ring ring-name {
    data-channel {
      vlan number
    }
    east-interface {
      control-channel channel-name {
        vlan number;
        interface name interface-name
      }
    }
    guard-interval number;
    node-id mac-address;
    restore-interval number;
    ring-protection-link-owner;
    non-revertive;
    wait-to-block-interval number;
    major-ring-name name;
    propagate-tc;
    compatibility-version (1|2);
    ring-id number;
    non-vc-mode;
    dot1p-priority number;
    west-interface {
      control-channel channel-name {
        vlan number;
        interface name interface-name
      }
      virtual-control-channel {
        west-interface name;
        east-interface name;
      }
    }
  }
 control-vlan (vlan-id | vlan-name);
    east-interface {
      node-id mac-address;
      control-channel channel-name {
        vlan number;
        interface name interface-name
      }
      interface-none
      ring-protection-link-end;
    }
  }
    control-channel channel-name {
      vlan number;
      interface name interface-name
    }
  }
  data-channel {
    vlan number
```

Copyright © 2018, Juniper Networks, Inc.

```
    }
    guard-interval number;
    node-id mac-address;
    restore-interval number;
    ring-protection-link-owner;
      west-interface {
        node-id mac-address;
        control-channel channel-name {
          vlan number;
          interface name interface-name
        }
        interface-none
        ring-protection-link-end;
      }
      control-channel channel-name {
        vlan number;
        interface name interface-name
      }
    }
  }
  guard-interval number;
  restore-interval number;
  traceoptions {
    file filename <no-stamp> <world-readable | no-world-readable> <replace> <size size>;
    flag flag;
  }
}
```

Hierarchy Level [edit protocols]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description Configure Ethernet ring protection switching.

The statements are explained separately. All statements apply to MX Series routers. EX Series switches do not assign **node-id** and use **control-vlan** instead of **control-channel**.

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

Related Documentation

- *Ethernet Ring Protection Switching Overview*
- *Ethernet Ring Protection Using Ring Instances for Load Balancing*
- *Example: Configuring Load Balancing Within Ethernet Ring Protection for MX Series Routers*
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*

 Copyright © 2018, Juniper Networks, Inc.

protocol-down

Syntax protocol-down;

Hierarchy Level [edit protocols oam ethernet link-fault-management action-profile event]

Release Information Statement introduced in Junos OS Release 8.5.

Description Upper layer indication of protocol down event. When the **protocol-down** statement is included, the protocol down event triggers the action specified under the **action** statement.

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

Related Documentation

- *Configuring an OAM Action Profile*

protocols

Syntax protocols [inet iso mpls];

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

Release Information Statement introduced in Junos OS Release 8.3.

Description For Layer 2.5 VPNs on T Series, MX Series, M120, and M320 routers support, configure IS-IS (ISO traffic) or MPLS traffic to traverse a TCC interface. By default, IPv4 (inet) traffic runs on T Series, MX, Series, M120, and M320 routers and over TCC interfaces. You must configure the same traffic type on both ends of the Layer 2.5 VPN.

NOTE: Some platform and FPC combinations can not pass TCC encapsulated ISO traffic. See *Platforms/FPCs That Cannot Forward TCC Encapsulated ISO Traffic* for details.

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

Related Documentation

- Configuring IS-IS or MPLS Traffic for TCC Interfaces on page 281
- *Platforms/FPCs That Cannot Forward TCC Encapsulated ISO Traffic*

proxy

Syntax proxy inet-address *address*;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description For Layer 2.5 VPNs using an Ethernet interface as the TCC router, configure the IP address for which the TCC router is proxying. Ethernet TCC is supported on interfaces that carry IPv4 traffic only. Ethernet TCC encapsulation is supported on 1-port Gigabit Ethernet, 2-port Gigabit Ethernet, 4-port Gigabit Ethernet, and 4-port Fast Ethernet PICs only. Ethernet TCC is not supported on the T640 router.

Options **inet-address**—Configure the IP address of the neighbor to the TCC router.

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

Related Documentation

- *Configuring Translation Cross-Connect Interface Switching*
- *Example: Configuring an Ethernet TCC or Extended VLAN TCC*
- remote on page 926
- *Junos OS VPNs Library for Routing Devices*

 Copyright © 2018, Juniper Networks, Inc.

proxy-arp

Syntax proxy-arp (restricted | unrestricted);

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

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

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

NOTE: You must configure the IP address and the **inet** family for the interface when you enable proxy ARP.

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

Options

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

Default: unrestricted

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

Related Documentation

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

push

Syntax

push;

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.

Description

NOTE: On EX4300 switches, **push** is not supported at the [edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map] hierarchy level.

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

You can use this statement on Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces; 10-Gigabit Ethernet LAN/WAN PIC; aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces; 100-Gigabit Ethernet Type 5 PIC with CFP; and Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces.

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

Required Privilege Level

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

Related Documentation

- *Stacking a VLAN Tag*
- *Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support (CLI Procedure)*

Copyright © 2018, Juniper Networks, Inc.

push-push

Syntax push-push;

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

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify the VLAN rewrite operation to push two VLAN tags in front of the frame.

You can use this statement on Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP.

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

Related Documentation

- *Stacking Two VLAN Tags*

queue-depth

Syntax queue-depth *cells*;

Hierarchy Level [edit interfaces *interface-name* atm-options linear-red-profiles *profile-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For ATM2 IQ interfaces only, define maximum queue depth in the CoS VC drop profile. Packets are always dropped beyond the defined maximum. This statement is mandatory; there is no default configuration.

Default Buffer usage is unregulated.

Options *cells*—Maximum number of cells the queue can contain.

Range: 1 through 64,000 cells

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

Related Documentation

- *ATM2 IQ VC Tunnel CoS Components Overview*
- *Configuring Linear RED Profiles on ATM Interfaces*
- high-plp-threshold on page 636
- low-plp-threshold on page 759

 Copyright © 2018, Juniper Networks, Inc.

queue-length

Syntax queue-length *number*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* address *address* family *family* multipoint-destination *address* shaping],
[edit interfaces *interface-name* unit *logical-unit-number* shaping],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* address *address* family *family* multipoint-destination *address* shaping],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* shaping]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description For ATM1 interfaces only, define the maximum queue length in the traffic-shaping profile. For ATM1 PICs, each VC has its own independent shaping parameters.

Default Buffer usage is unregulated.

Options *number*—Maximum number of packets the queue can contain.

Range: 1 through 16,383 packets

Default: 16,383 packets

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

Related Documentation
- *Configuring the ATM1 Queue Length*

Copyright © 2018, Juniper Networks, Inc.

queues

Syntax queues [*queue-numbers*];

Hierarchy Level [edit interfaces vsp-*fpc/pic/port* unit *logical-unit-number* compression rtp]

Release Information Statement introduced before Junos OS Release 7.4.

Description For voice services interfaces only, assign queue numbers for RTP traffic.

Options **queues *queue-numbers***—Assign one or more of the following queues: **q0**, **q1**, **q2**, **q3**. For VRRP services, specify the **q3** option instead of **q0**.

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

quiet-period

Syntax quiet-period *seconds*;

Hierarchy Level [edit protocols dot1x authenticator interface *interface-id*]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the number of seconds the port remains in the wait state following a failed authentication exchange with the client, before reattempting authentication.

Options ***seconds***—Specify the number of seconds the port remains in the wait state following a failed authentication exchange with the client, before reattempting authentication.

Range: 0 through 65,535 seconds

Default: 60 seconds

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

Related Documentation

- *IEEE 802.1x Port-Based Network Access Control Overview*
- authenticator on page 448
- dot1x on page 543
- interface (IEEE 802.1x) on page 682

 Copyright © 2018, Juniper Networks, Inc.

radius-realm

Syntax radius-realm *radius-realm-string*;

Hierarchy Level [edit interfaces *interface-name* auto-configure vlan-ranges authentication username-include],
[edit interfaces *interface-name* auto-configure stacked-vlan-ranges authentication username-include]

Release Information Statement introduced in Junos OS Release 10.0.

Description Specify that the user-defined RADIUS realm string is appended as a last piece to the username and used by RADIUS to direct the authentication request to a profile that does not allocates addresses.

Options ***radius-realm-string***—A string to describe the RADIUS realm.

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

Related Documentation
- *Configuring VLAN Interface Username Information for AAA Authentication*

ranges (Dynamic Stacked VLAN)

Syntax ranges (any | *low-tag*–*high-tag*),(any | *low-tag*–*high-tag*);

Hierarchy Level [edit interfaces *interface-name* auto-configure stacked-vlan-ranges dynamic-profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure VLAN ranges for dynamic, auto-sensed stacked VLANs.

Options **any**—The entire VLAN range.

low-tag—The lower limit of the VLAN range.

high-tag—The upper limit of the VLAN range.
Range: 1 through 4094

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

Related Documentation
- *Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs*

Copyright © 2018, Juniper Networks, Inc.

ranges (Dynamic VLAN)

Syntax ranges (any | *low-tag*)-(any | *high-tag*);

Hierarchy Level [edit interfaces *interface-name* auto-configure vlan-ranges dynamic-profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure VLAN ranges for dynamic, auto-sensed VLANs.

Options **any**—The entire VLAN range.

low-tag—The lower limit of the VLAN range.

high-tag—The upper limit of the VLAN range.

Range: 1 through 4094

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

Related Documentation

- *Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs*

rate

Syntax rate *percentage*;

Hierarchy Level [edit interfaces *interface-name* receive-bucket],
[edit interfaces *interface-name* transmit-bucket]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify percentage of the interface line rate that is available to receive or transmit packets.

Options ***percentage***—Percentage of the interface line rate that is available to receive or transmit packets.

Range: 0 through 100

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

Related Documentation

- Configuring Receive and Transmit Leaky Bucket Properties to Reduce Network Congestion on page 156

 Copyright © 2018, Juniper Networks, Inc.

rate

Syntax rate *new-sessions-per-second*;

Hierarchy Level [edit interfaces *interface-name* services-options session-limit]

Release Information Statement introduced in Junos OS Release 9.6.

Description Specify the maximum number of new sessions allowed per second.

Options **rate *new-sessions-per-second***—Specify the maximum number of new sessions allowed per second.

Range: 0, which indicates no limit, and 500 or greater.

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

rdi

Syntax rdi;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management action-profile *tlv-action* event]

Release Information Statement introduced in Junos OS Release 10.1.

Description Define a new event **rdi**. The remote defect indication (rdi) event is triggered whenever CCM packets are received from a remote location with the rdi bit set.

This event is cleared and action is reverted when none of the remote MEPs send the CCM packets with the RDI bit.

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

Related Documentation

- *Configuring a CFM Action Profile to Specify CFM Actions for CFM Events*

Copyright © 2018, Juniper Networks, Inc.

reassemble-packets

Syntax reassemble-packets;

Hierarchy Level [edit interfaces gr-*fpc/pic/port* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces gr-*fpc/pic/port* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Enable reassembly of fragmented tunnel packets on generic routing encapsulation (GRE) tunnel interfaces.

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*
- *Configuring Packet Reassembly*

reauthentication

Syntax reauthentication (disable | interval *seconds*);

Hierarchy Level [edit protocols dot1x authenticator interface *interface-id*]

Release Information Statement introduced in Junos OS Release 9.3.

Description Set or disable the periodic reauthentication of the client.

Options

- **disable**—Disable the periodic reauthentication of the client.
- **interval *seconds***—Specify the periodic reauthentication time interval.

Range: 1 through 65,535 seconds
Default: 3600 seconds

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

Related Documentation

- *IEEE 802.1x Port-Based Network Access Control Overview*
- dot1x on page 543
- interface (IEEE 802.1x) on page 682
- quiet-period on page 916

 Copyright © 2018, Juniper Networks, Inc.

receive-bucket

Syntax

```
receive-bucket {
  overflow (discard | tag);
  rate percentage;
  threshold bytes;
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Set parameters for the receive leaky bucket, which specifies what percentage of the interface's total capacity can be used to receive packets.

For each DS3 channel on a channelized OC12 interface, you can configure a unique receive bucket.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- Configuring Receive and Transmit Leaky Bucket Properties to Reduce Network Congestion on page 156
- **transmit-bucket on page 1040**

receive-options-packets

Syntax receive-options-packets;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description For a Monitoring Services PIC and an ATM or SONET/SDH PIC installed in an M160, M40e, or T Series router, guarantee conformity with cflowd records structure. This statement is required when you enable passive monitoring.

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

Related Documentation

- *Enabling Passive Monitoring on ATM Interfaces*
- *Enabling Packet Flow Monitoring on SONET/SDH Interfaces*

receive-ttl-exceeded

Syntax receive-ttl-exceeded;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description For Monitoring Services PIC and an ATM or SONET/SDH PIC installed in an M160, M40e, or T Series router, guarantee conformity with cflowd records structure. This statement is required when you enable passive monitoring.

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

Related Documentation

- *Enabling Passive Monitoring on ATM Interfaces*
- *Enabling Packet Flow Monitoring on SONET/SDH Interfaces*

 Copyright © 2018, Juniper Networks, Inc.

red-differential-delay

Syntax red-differential-delay *milliseconds*;

Hierarchy Level [edit interfaces *interface-name* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services and voice services interfaces only, configure the red differential delay among bundle links to give warning when a link has a differential delay that exceeds the configured threshold.

Options ***milliseconds***—Red differential delay threshold.

Range: 1 through 2000 milliseconds

Default: 10 milliseconds

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

Related Documentation

- action-red-differential-delay on page 421
- yellow-differential-delay on page 1106
- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2018, Juniper Networks, Inc.

redial-delay

Syntax

redial-delay *time*;

Hierarchy Level

[edit interfaces dl*n* unit *logical-unit-number* dialer-options],
[edit logical-systems *logical-system-name* interfaces dl*n* unit *logical-unit-number* dialer-options]

Release Information

Statement introduced in Junos OS Release 7.5.

Description

On J Series Services Routers with interfaces configured for ISDN with dialout, specify the delay (in seconds) between two successive calls made by the dialer. To configure callback mode, include the **callback** statement at the **[edit interfaces dl*n* unit *logical-unit-number* dialer-options]** hierarchy level.

If the **callback** statement is configured, you cannot use the **caller *caller-id*** statement at the **[edit interfaces dl*n* unit *logical-unit-number* dialer-options]** hierarchy level.

Options

time—Delay (in seconds) between two successive calls.

Range: 2 through 255 seconds

Default: 3 seconds

Required Privilege Level

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

Related Documentation

- *ISDN Interfaces Overview*
- *Junos OS Interfaces and Routing Configuration Guide*

 Copyright © 2018, Juniper Networks, Inc.

redundancy-options

Syntax

```
redundancy-options {
  primary interface-name;
  secondary interface-name;
  hot-standby;
}
```

Hierarchy Level

[edit interfaces (rsp0 | rsp1)],
[edit interfaces rlsq*number*]
[edit interfaces rsp*number*]
[edit interfaces rms*number*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the primary and secondary (backup) AS PIC interfaces or MultiServices PIC interfaces.

Options

primary *interface-name*—The identifier for the primary LSQ AS, rsp, or rms interface.

secondary *interface-name*—The identifier for the secondary (backup) LSQ AS, rsp, or rmsinterface.

hot-standby—For one-to-one AS, rsp, or rms redundancy configurations, specify that the failure detection and recovery must take place in less than 5 seconds.

Required Privilege Level

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

remote

Syntax

```
remote {
  (inet-address address | mac-address address);
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

For Layer 2.5 VPNs using an Ethernet interface as the TCC router, configure the location of the remote router. Ethernet TCC is supported on interfaces that carry IPv4 traffic only. Ethernet TCC encapsulation is supported on 1-port Gigabit Ethernet, 2-port Gigabit Ethernet, 4-port Gigabit Ethernet, and 4-port Fast Ethernet PICs only.

Options

mac-address—Configure the MAC address of the remote site.

inet-address—Configure the IP address of the remote site.

Required Privilege Level

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

Related Documentation

- *Configuring Translation Cross-Connect Interface Switching*
- *Example: Configuring an Ethernet TCC or Extended VLAN TCC*
- proxy on page 910
- *Junos OS VPNs Library for Routing Devices*

 Copyright © 2018, Juniper Networks, Inc.

remote-loopback

Syntax remote-loopback;

Hierarchy Level [edit protocols oam link-fault-management interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.2.

Description For Ethernet interfaces on EX Series switches and M320, M120, MX Series, and T Series routers, set the remote DTE into loopback mode. Remove the statement from the configuration to take the remote DTE out of loopback mode. Used for IEEE 802.3ah Operation, Administration, and Management (OAM) support.

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

Related Documentation

- *Setting a Remote Interface into Loopback Mode*

remote-loopback-respond

Syntax remote-loopback-respond;

Hierarchy Level [edit interfaces ct1-*fpc/pic/port*],
[edit interfaces *interface-name* t1-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers.

Description For T1 interfaces only, configure the router to respond to remote loopback requests. Remote loopback requests can be from the facilities data link or inband.

NOTE: When configuring CT1 interfaces on the 10-port Channelized E1/T1 IQE PIC, the **remote-loopback-respond** statement must be included at the **[edit interfaces ct1-*fpc/pic/port*]** hierarchy level.

Default The router does not respond to remote loop requests.

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

Related Documentation

- *Configuring the T1 Remote Loopback Response*
- feac-loop-respond on page 614
- loopback (ADSL, DS0, E1/E3, SONET/SDH, SHDSL, and T1/T3) on page 751

 Copyright © 2018, Juniper Networks, Inc.

remote-mep

Syntax

```
remote-mep mep-id {
  action-profile profile-name;
  sla-iterator-profile profile-name {
    data-tlv-size size;
    iteration-count count-value;
    priority priority-value;
  }
  detect-loc;
}
```

Hierarchy Level

[edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id*]

Release Information

Statement introduced in Junos OS Release 8.4.

Description

Configure the numeric identifier of the remote maintenance association end point (MEP) within the maintenance association.

Options

mep-id—Numeric identifier of the MEP.

Range: 1 through 8191

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- *Configuring a MEP to Generate and Respond to CFM Protocol Messages*
- *detect-loc*

remove-when-no-subscribers

Syntax remove-when-no-subscribers;

Hierarchy Level [edit interfaces *interface-name* auto-configure]

Release Information Statement introduced in Junos OS Release 11.4.

Description Remove subscriber VLANs automatically when no client sessions (for example, DHCP or PPPoE) exist on the VLAN.

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

Related Documentation
- *Automatically Removing VLANs with No Subscribers*

request

Syntax request (protect | working);

Hierarchy Level [edit interfaces *interface-name* sonet-options aps]

Release Information Statement introduced before Junos OS Release 7.4.

Description Perform a manual switch between the protect and working circuits. This statement is honored only if there are no higher-priority reasons to switch.

Options **protect**—Request that the circuit become the protect circuit.

working—Request that the circuit become the working circuit.

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

Related Documentation
- *Configuring Switching Between the Working and Protect Circuits*
- force on page 621

 Copyright © 2018, Juniper Networks, Inc.

required-depth

Syntax

required-depth *number*;

Hierarchy Level

[edit interfaces *interface-name* atm-options mpls pop-all-labels],
[edit interfaces *interface-name* sonet-options mpls pop-all-labels],
[edit interfaces *interface-name* fastether-options mpls pop-all-labels],
[edit interfaces *interface-name* gigether-options mpls pop-all-labels]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description

For passive monitoring on ATM and SONET/SDH interfaces only, specify the number of MPLS labels an incoming packet must have for the **pop-all-labels** statement to take effect.

If you include the **required-depth 1** statement, the **pop-all-labels** statement takes effect for incoming packets with one label only. If you include the **required-depth 2** statement, the **pop-all-labels** statement takes effect for incoming packets with two labels only.

Options

number—Number of MPLS labels on incoming IP packets.

Range: 1 or 2 labels

Default: If you omit this statement, the **pop-all-labels** statement takes effect for incoming packets with one or two labels. The default is equivalent to including the **required-depth [1 2]** statement.

Required Privilege Level

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

Related Documentation

- *Removing MPLS Labels from Incoming Packets*
- *Enabling Packet Flow Monitoring on SONET/SDH Interfaces*
- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2018, Juniper Networks, Inc.

restore-interval

Syntax restore-interval *number*;

Hierarchy Level [edit protocols protection-group ethernet-ring *ring-name*]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

Description Configures the number of minutes that the node does not process any Ethernet ring protection (ERP) protocol data units (PDUs).. This configuration is a global configuration and applies to all Ethernet rings if the Ethernet ring does not have a more specific configuration for this value. If no parameter is configured at the protection group level, the global configuration of this parameter uses the default value.

Options *number*—Specify the restore interval.

Range: 1 through 12 minutes

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

Related Documentation

- *Ethernet Ring Protection Switching Overview*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

 Copyright © 2018, Juniper Networks, Inc.

retries

Syntax retries *integer*;

Hierarchy Level [edit protocols dot1x authenticator interface *interface-id*]

Release Information Statement introduced in Junos OS Release 9.3.

Description Set a limit on the number of failed authentication attempts between a port and a client. When the limit is exceeded, the port waits to reattempt authentication for the number of seconds set by the **quiet-period** statement configured at the same hierarchy level.

Options *integer*—Specify the number of retries.

Range: 1 through 10

Default: 3 retries

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

Related Documentation

- *IEEE 802.1x Port-Based Network Access Control Overview*
- dot1x on page 543
- interface (IEEE 802.1x) on page 682
- quiet-period on page 916

Copyright © 2018, Juniper Networks, Inc.

revert-time (Interfaces)

Syntax revert-time *seconds*;

Hierarchy Level [edit interfaces *interface-name* sonet-options aps]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure APS revertive mode.

Default APS operates in nonrevertive mode.

Options ***seconds***—Amount of time to wait after the working circuit has again become functional before making the working circuit active again.

Range: 1 through 65,535 seconds

Default: None (APS operates in nonrevertive mode)

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

Related Documentation
- *Configuring Revertive Mode*

 Copyright © 2018, Juniper Networks, Inc.

revertive

Syntax revertive;

Hierarchy Level [edit interfaces ae*X* aggregated-ether-options lacp link-protection]

Release Information Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 12.3 for EX Series switches.
Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.

Description Enable the ability to switch to a better priority link (if one is available).

NOTE: By default, LACP link protection is revertive. However, you can use this statement to define a specific aggregated Ethernet interface as revertive to override a global non-revertive statement specified at the [edit chassis] hierarchy level.

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

Related Documentation
- *non-revertive (Chassis)*
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*

rfc-2615

Syntax rfc-2615;

Hierarchy Level [edit interfaces *interface-name* sonet-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Include this statement to enable features described in RFC 2615, *PPP over SONET/SDH*.

Default Settings required by RFC 1619, *PPP over SONET/SDH*.

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

Related Documentation
- *Configuring PPP Support on SONET/SDH Interfaces*

Copyright © 2018, Juniper Networks, Inc.

ring-protection-link-end

Syntax ring-protection-link-end;

Hierarchy Level [edit protocols protection-group ethernet-ring *ring-name* (east-interface | west-interface)]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description Specify that the port is one side of a ring protection link (RPL) by setting the RPL end flag.

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

Related Documentation

- *Ethernet Ring Protection Switching Overview*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

ring-protection-link-owner

Syntax ring-protection-link-owner;

Hierarchy Level [edit protocols protection-group ethernet-ring *ring-name*]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

Description Specify the ring protection link (RPL) owner flag in the Ethernet protection ring. Include this statement only once for each ring (only one node can function as the RPL owner).

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

Related Documentation

- *Ethernet Ring Protection Switching Overview*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*

Copyright © 2018, Juniper Networks, Inc.

routing-instance

Syntax

```
routing-instance {
  destination routing-instance-name;
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

To configure **interfaces** and **logical-systems**, specify the destination routing instance that points to the routing table containing the tunnel destination address.

Default

The default Internet routing table is **inet.0**.

Required Privilege Level

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2018, Juniper Networks, Inc.

routing-instance (PPPoE Service Name Tables)

Syntax routing-instance *routing-instance-name*;

Hierarchy Level [edit protocols pppoe service-name-tables *table-name* service *service-name*],
[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Use in conjunction with the **dynamic-profile** statement at the same hierarchy levels to specify the routing instance in which to instantiate a dynamic PPPoE interface. You can associate a routing instance with a named service entry, **empty** service entry, or **any** service entry configured in a PPPoE service name table, or with an agent circuit identifier/agent remote identifier (ACI/ARI) pair defined for these services.

The routing instance associated with a service entry in a PPPoE service name table overrides the routing instance associated with the PPPoE underlying interface on which the dynamic PPPoE interface is created.

If you include the **routing-instance** statement at the **[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]** hierarchy level, you cannot also include the **static-interface** statement at this level. The **routing-instance** and **static-interface** statements are mutually exclusive for ACI/ARI pair configurations.

Options *routing-instance-name*—Name of the routing instance in which the router instantiates the dynamic PPPoE interface.

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

Related Documentation

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

 Copyright © 2018, Juniper Networks, Inc.

rpf-check (Dynamic Profiles)

Syntax

```
rpf-check {
  fail-filter filter-name;
  mode loose;
}
```

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

Release Information Statement introduced in Junos OS Release 9.6.

Description Check whether traffic is arriving on an expected path. You can include this statement with the **inet** protocol family only.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

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

Related Documentation

- Configuring Unicast RPF Strict Mode on page 244
- *Configuring Unicast RPF and Fail Filters in Dynamic Profiles for Subscriber Interfaces*

Copyright © 2018, Juniper Networks, Inc.

rpf-check

List of Syntax

Syntax (MX Series, SRX Series, M Series, T Series, PTX Series) on page 940
Syntax (EX Series) on page 940

Syntax (MX Series, SRX Series, M Series, T Series, PTX Series)

```
rpf-check {
  fail-filter filter-name;
  mode loose;
}
```

Syntax (EX Series)

```
rpf-check;
```

Hierarchy Level (MX Series, SRX Series, M Series, T Series, PTX Series)

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

Hierarchy Level (EX Series)

[edit interfaces *interface-name* unit *logical-unit-number* family inet],
[edit interfaces *interface-name* unit *logical-unit-number* family inet6]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Support for interface **ps0** (pseudowire subscriber logical interface device) added in Junos OS Release 15.1.

Description

Enable a reverse-path forwarding (RPF) check on unicast traffic.

On EX3200 and EX4200 switches, enable a reverse-path forwarding (RPF) check on unicast traffic (except ECMP packets) on all ingress interfaces.

On EX4300 switches, enable a reverse-path forwarding (RPF) check on unicast traffic, including ECMP packets, on all ingress interfaces.

On EX8200 and EX6200 switches, enable an RPF check on unicast traffic, including ECMP packets, on the selected ingress interfaces.

On QFX Series switches, enable an RPF check on unicast traffic (except ECMP packets) on the selected ingress interfaces.

The mode statement is explained separately.

Default

Unicast RPF is disabled on all interfaces.

 Copyright © 2018, Juniper Networks, Inc.

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

- Configuring Unicast RPF Strict Mode on page 244
- Configuring Unicast RPF Loose Mode on page 246
- Example: Configuring Unicast Reverse-Path-Forwarding Check on page 253
- *Configuring a Pseudowire Subscriber Logical Interface Device*
- *Example: Configuring Unicast RPF on an EX Series Switch*
- *Configuring Unicast RPF (CLI Procedure)*
- *Disabling Unicast RPF (CLI Procedure)*
- *Understanding Unicast RPF*

rpf-loose-mode-discard

Syntax

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

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

- Configuring Unicast RPF on page 241

Copyright © 2018, Juniper Networks, Inc.

rtp

Syntax

```
rtp {
  f-max-period number;
  queues [ queue-numbers ];
  port {
    minimum port-number;
    maximum port-number;
  }
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* compression]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the real-time transport protocol (RTP) properties for voice services traffic.

The remaining statements are described separately.

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

 Copyright © 2018, Juniper Networks, Inc.

rts

Syntax rts (assert | de-assert | normal);

Hierarchy Level [edit interfaces *interface-name* serial-options dce-options],
[edit interfaces *interface-name* serial-options dte-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For EIA-530 and V.35 interfaces only, configure the to-DCE signal, request to send (RTS).

Options **assert**—The to-DCE signal must be asserted.

de-assert—The to-DCE signal must be deasserted.

normal—Normal RTS signal handling, as defined by the TIA/EIA Standard 530.

Default: normal

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

Related Documentation

- Configuring the Serial Signal Handling on page 344

rts-polarity

Syntax rts-polarity (negative | positive);

Hierarchy Level [edit interfaces *interface-name* serial-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure RTS signal polarity.

Options **negative**—Negative signal polarity.

positive—Positive signal polarity.

Default: positive

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

Related Documentation

- Configuring Serial Signal Polarities on page 347

Copyright © 2018, Juniper Networks, Inc.

rtvbr

Syntax

rtvbr peak *rate* sustained *rate* burst *length*;

Hierarchy Level

[edit interfaces *interface-name* atm-options vpi *vpi-identifier* shaping],
[edit interfaces *interface-name* unit *logical-unit-number* address *address* family *family* multipoint-destination *address* shaping],
[edit interfaces *interface-name* unit *logical-unit-number* shaping],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* shaping],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* address *address* family *family* multipoint-destination *address* shaping]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

For ATM2 IQ PICs only, define the real-time variable bandwidth utilization in the traffic-shaping profile.

When you configure the real-time bandwidth utilization, you must specify all three options (**burst**, **peak**, and **sustained**). You can specify the rate 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 the rate 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 using the formula 1 cps = 384 bps.

Default

If the **rtvbr** statement is not included, bandwidth utilization is unlimited.

Options

burst ***length***—Burst length, in cells. If you set the length to 1, the peak traffic rate is used.

Range: 1 through 4000 cells

peak ***rate***—Peak rate, in bits per second or cells per second.

Range: For ATM2 IQ OC3 and OC12 interfaces, 33 Kbps through 542,526,792 bps. For ATM2 IQ OC48 interfaces, 33 Kbps through 2,170,107,168 bps. For ATM2 IQ DS3 and E3 interfaces, 33 Kbps through the maximum rate, which depends on the ATM encapsulation and framing you configure..

sustained ***rate***—Sustained rate, in bps or cps.

Range: For ATM2 IQ OC3 and OC12 interfaces, 33 Kbps through 542,526,792 bps. For ATM2 IQ OC48 interfaces, 33 Kbps through 2,170,107,168 bps. For ATM2 IQ DS3 and E3 interfaces, from 33 Kbps through the maximum rate, which depends on the ATM encapsulation and framing you configure.

Required Privilege Level

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

Related Documentation

- *Configuring ATM CBR*

Copyright © 2018, Juniper Networks, Inc.

- *Configuring ATM2 IQ Real-Time VBR*
- *Applying Scheduler Maps to Logical ATM Interfaces*
- cbr on page 478
- vbr on page 1071

Copyright © 2018, Juniper Networks, Inc.

sa-multicast (100-Gigabit Ethernet)

Syntax sa-multicast;

Hierarchy Level [edit chassis fpc *slot* pic *slot* forwarding-mode]

Release Information Statement introduced in Junos OS Release 10.4.

Description Configure the 100-Gigabit Ethernet PIC or MIC to interoperate with other Juniper Networks 100-Gigabit Ethernet PICs.

NOTE: **The default packet steering mode for PD-1CE-CFP-FPC4 is SA multicast bit mode. No SA multicast configuration is required to enable this mode.**

sa-multicast supports interoperability between the following PICs and MICs:

- 100-Gigabit Ethernet Type 5 PIC with CFP (PF-1CGE-CFP) and the 100-Gigabit Ethernet Type 4 PIC with CFP (PD-1CE-CFP-FPC4) .
- 100-Gigabit Ethernet MICs and the 100-Gigabit Ethernet Type 4 PIC with CFP (PD-1CE-CFP-FPC4).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP*
- *Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4*
- *Configuring 100-Gigabit Ethernet MICs to Interoperate with Type 4 100-Gigabit Ethernet PICs (PD-1CE-CFP-FPC4) Using SA Multicast Mode*
- *Interoperability Between MPC4E (MPC4E-3D-2CGE-8XGE) and 100-Gigabit Ethernet PICs on Type 4 FPC*
- *Configuring MPC4E (MPC4E-3D-2CGE-8XGE) to Interoperate with 100-Gigabit Ethernet PICs on Type 4 FPC Using SA Multicast Mode*
- *Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP*
- *Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4*
- *forwarding-mode (100-Gigabit Ethernet)*

 Copyright © 2018, Juniper Networks, Inc.

- *sa-multicast (PTX Series Packet Transport Routers)*
- vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP) on page 1088
- *Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP*

sampling (Interfaces)

Syntax sampling *direction*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet]

Release Information Statement introduced before Junos OS Release 7.4.
Support for **sampling** on both input and output for bridge family introduced in Junos OS Release 18.2R1.
Support for **sampling** on both input and output for vpls family introduced in Junos OS Release 18.2R1.

Description Configure the direction of traffic to be sampled.

Options ***direction*** can be one of the following:

input—Configure at least one expected ingress point.

output—Configure at least one expected egress point.

input output—On a single interface, configure at least one expected ingress point and one expect egress point.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*
- *Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches*

Copyright © 2018, Juniper Networks, Inc.

satop-options

Syntax

```
satop-options {
  excessive-packet-loss-rate {
    apply-groups group-name
    apply-groups-except group-name
    groups group-name
    sample-period milliseconds
    threshold percentile
  )
  idle-pattern pattern
  jitter-buffer-auto-adjust
  jitter-buffer-latency milliseconds
  jitter-buffer-packets packets
  payload-size bytes;
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers.

Description Set Structure-Agnostic TDM over Packet (SAToP) protocol options.

On ACX Series routers, the following statements are not supported:

```
apply-groups group-name
apply-groups-except group-name
groups group-name
jitter-buffer-auto-adjust
```

Options **excessive-packet-loss-rate options**—Set packet loss options.

- **apply-groups *group-name***—Groups from which to inherit configuration data.
- **apply-groups-except *group-name***—Don't inherit configuration data from these groups.
- **groups *group-name***—Specify groups.
- **sample-period *milliseconds***—Number of milliseconds over which excessive packet loss rate is calculated.
- **threshold *percentile***—Percentile designating the threshold of excessive packet loss rate (from 1 to 100).

idle-pattern *pattern*—An 8–bit hexadecimal pattern to replace TDM data in a lost packet (from 0 to 255).

jitter-buffer-auto-adjust—Automatically adjust the jitter buffer.

NOTE: This option is not applicable on MX Series routers.

Copyright © 2018, Juniper Networks, Inc.

jitter-buffer-latency *milliseconds*—Number of milliseconds delay in jitter buffer (from 1 to 1000 milliseconds).

jitter-buffer-packets *packets*—Number of packets in jitter buffer (from 1 to 64).

payload-size *bytes*—Payload size in integer number of bytes.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring SAToP on 4-Port Channelized OC3/STM1 Circuit Emulation MICs*
- *Configuring SAToP Emulation on T1/E1 Interfaces on 12-Port Channelized T1/E1 Circuit Emulation PICs*
- *ATM Support on Circuit Emulation PICs Overview*

shared-interface

Syntax

shared-interface;

Hierarchy Level

[edit interfaces ge-*fpc/pic/slot*],
[edit interfaces so-*fpc/pic/slot*],
[edit interfaces xe-*fpc/pic/slot*]

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Configure a physical interface to be a shared interface. Logical interfaces configured under the shared physical interface can be assigned to different Protected System Domains (PSDs).

Options

This statement has no options.

Required Privilege Level

view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation

- interface-shared-with on page 691

scheduler-maps (For ATM2 IQ Interfaces)

Syntax

scheduler-maps *map-name* {
  forwarding-class (*class-name* | assured-forwarding | best-effort | expedited-forwarding
    | network-control);
  vc-cos-mode (alternate | strict);
}

Hierarchy Level

[edit at-*fpc/pic/port interface-name* atm-options]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

For ATM2 IQ interfaces only, define CoS parameters assigned to forwarding classes.

Options

map-name—Name of the scheduler map.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *ATM2 IQ VC Tunnel CoS Components Overview*
- *Applying Scheduler Maps to ATM Interfaces*
- atm-scheduler-map on page 445

 Copyright © 2018, Juniper Networks, Inc.

schedulers

Syntax schedulers *number*;

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 8.2.

Description Specify the number of schedulers for Ethernet IQ2 and IQ2-E PIC port interfaces.

Default If you omit this statement, the 1024 schedulers are distributed equally over all ports in multiples of 4.

Options ***number***—Number of schedulers to configure on the port.

Range: 1 through 1024

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Class of Service Feature Guide for Routing Devices and EX9200 Switches*

secondary

Syntax secondary *interface-name*;

Hierarchy Level [edit interfaces (rsp0 | rsp1) redundancy-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the secondary (backup) AS PIC interface or MultiServices PIC interface.

Options ***interface-name***—The identifier for the AS PIC interface or MultiServices PIC interface, which must be of the form **sp-*fpc/pic/port***.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2018, Juniper Networks, Inc.

send-critical-event

Syntax send-critical-event;

Hierarchy Level [edit protocols oam ethernet link-fault-management action-profile action]

Release Information Statement introduced in Junos OS Release 8.5.

Description Send OAM PDUs with the critical event bit set.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Specifying the Actions to Be Taken for Link-Fault Management Events*

Copyright © 2018, Juniper Networks, Inc.

serial-options

Syntax

```
serial-options {
  clock-rate rate;
  clocking-mode (dce | loop);
  control-polarity (negative | positive);
  cts-polarity (negative | positive);
  dcd-polarity (negative | positive);
  dce-options {
    control-signal (assert | de-assert | normal);
    cts (ignore | normal | require);
    dcd (ignore | normal | require);
    dsr (ignore | normal | require);
    dtr signal-handling-option;
    ignore-all;
    indication (ignore | normal | require);
    rts (assert | de-assert | normal);
    tm (ignore | normal | require);
  }
  dsr-polarity (negative | positive);
  dte-options {
    control-signal (assert | de-assert | normal);
    cts (ignore | normal | require);
    dcd (ignore | normal | require);
    dsr (ignore | normal | require);
    dtr signal-handling-option;
    ignore-all;
    indication (ignore | normal | require);
    rts (assert | de-assert | normal);
    tm (ignore | normal | require);
  }
  dtr-circuit (balanced | unbalanced);
  dtr-polarity (negative | positive);
  encoding (nrz | nrzi);
  indication-polarity (negative | positive);
  line-protocol protocol;
  loopback (dce-local | dce-remote | local | remote);
  rts-polarity (negative | positive);
  tm-polarity (negative | positive);
  transmit-clock invert;
}
```

Hierarchy Level [edit interfaces se-*pim*/0/*port*]

Release Information Statement introduced prior to Junos OS Release 7.4.

Description Configure serial-specific interface properties.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Serial Interfaces Overview on page 335
- *no-concatenate*

server

Syntax server;

Hierarchy Level [edit interfaces pp0 unit *logical-unit-number* pppoe-options],
[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number* pppoe-options]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure the router to operate in the PPPoE server mode. Supported on M120 and M320 Multiservice Edge Routers and MX Series Universal Edge Routers operating as access concentrators.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring the PPPoE Server Mode*

 Copyright © 2018, Juniper Networks, Inc.

server-timeout

Syntax server-timeout *seconds*;

Hierarchy Level [edit protocols dot1x authenticator interface *interface-id*]

Release Information Statement introduced in Junos OS Release 9.3.

Description Sets the number of seconds the port waits for a reply when relaying a response from the client to the authentication server before timing out and invoking the server-fail action.

Options ***seconds***—The number of seconds the port waits for a response when relaying a request from the authentication server to the client before resending the request.

Range: 1 through 60 seconds

Default: 30 seconds

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *IEEE 802.1x Port-Based Network Access Control Overview*
- authenticator on page 448
- dot1x on page 543
- interface (IEEE 802.1x) on page 682

service (Logical Interfaces)

Syntax

```
service {
  input {
    service-set service-set-name <service-filter filter-name>;
    post-service-filter filter-name;
  }
  output {
    service-set service-set-name <service-filter filter-name>;
  }
}
```

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* family inet],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Define one or more service sets and filters, and one postservice filter to be applied to an interface.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2018, Juniper Networks, Inc.

service (PPPoE)

Syntax

```
service service-name {
  drop;
  delay seconds;
  terminate;
  dynamic-profile profile-name;
  routing-instance routing-instance-name;
  max-sessions number;
  agent-specifier {
    aci circuit-id-string ari remote-id-string {
      drop;
      delay seconds;
      terminate;
      dynamic-profile profile-name;
      routing-instance routing-instance-name;
      static-interface interface-name;
    }
  }
}
```

Hierarchy Level [edit protocols pppoe service-name-tables *table-name*]

Release Information Statement introduced in Junos OS Release 10.0.
any, **dynamic-profile**, **routing-instance**, **max-sessions**, and **static-interface** options introduced in Junos OS Release 10.2.

Description Specify the action taken by the interface on receipt of a PPPoE Active Discovery Initiation (PADI) control packet for the specified named service, **empty** service, or **any** service in a PPPoE service name table. You can also specify the dynamic profile and routing instance that the router uses to instantiate a dynamic PPPoE interface, and the maximum number of active PPPoE sessions that the router can establish with the specified service.

Default The default action is terminate.

Options *service-name*—Service entry in the PPPoE service name table:

- *service-name*—Named service entry of up to 32 characters; for example, **premiumService**. You can configure a maximum of 512 named service entries across all PPPoE service name tables on the router.
- **empty**—Service entry of zero length that represents an unspecified service. Each PPPoE service name table includes one **empty** service entry by default.
- **any**—Default service for non-empty service entries that do not match the named or **empty** service entries configured in the PPPoE service name table. Each PPPoE service name table includes one **any** service entry by default.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Copyright © 2018, Juniper Networks, Inc.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring PPPoE Service Name Tables*
- *Assigning a Service to a Service Name Table and Configuring the Action Taken When the Client Request Includes a Non-zero Service Name Tag*
- *Configuring the Action Taken When the Client Request Includes an Empty Service Name Tag*
- *Configuring the Action Taken for the Any Service*

service-domain

Syntax service-domain (inside | outside);

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For adaptive services interfaces, specify a service interface domain. If you specify this interface using the **next-hop-service** statement at the **[edit services service-set *service-set-name*]** hierarchy level, the interface domain must match that used with the **inside-service-interface** and **outside-service-interface** statements.

Options **inside**—Interface used within the network.

outside—Interface used outside the network.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

 Copyright © 2018, Juniper Networks, Inc.

service-filter (Interfaces)

Syntax service-filter *filter-name*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet service (input | output) service-set *service-set-name*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet service (input | output) service-set *service-set-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the filter to be applied to traffic before it is accepted for service processing. Configuration of a service filter is optional; if you include the **service-set** statement without a **service-filter** definition, Junos OS assumes the match condition is true and selects the service set for processing automatically.

Options *filter-name*—Identifies the filter to be applied in service processing. You can include special characters, such as a forward slash (/), colon (:), or a period (.).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Applying Filters and Services to Interfaces*
- *Junos OS Services Interfaces Library for Routing Devices*

service-name-table

Syntax service-name-table *table-name*;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family pppoe],
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]

Release Information Statement introduced in Junos OS Release 10.0.
Support at the **[edit ... family pppoe]** hierarchies introduced in Junos OS Release 11.2.

Description Specify the PPPoE service name table assigned to a PPPoE underlying interface. This underlying interface is configured with either the **encapsulation ppp-over-ether** statement or the **family pppoe** statement; the two statements are mutually exclusive.

NOTE: The **[edit ... family pppoe]** hierarchies are supported only on MX Series routers with MPCs.

Options *table-name*—Name of the PPPoE service name table, a string of up to 32 alphanumeric characters.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring PPPoE Service Name Tables*
- *Assigning a Service Name Table to a PPPoE Underlying Interface*
- *Configuring the PPPoE Family for an Underlying Interface*

 Copyright © 2018, Juniper Networks, Inc.

service-name-tables

Syntax

```
service-name-tables table-name {
  service service-name {
    drop;
    delay seconds;
    terminate;
    dynamic-profile profile-name;
    routing-instance routing-instance-name;
    max-sessions number;
    agent-specifier {
      aci circuit-id-string ari remote-id-string {
        drop;
        delay seconds;
        terminate;
        dynamic-profile profile-name;
        routing-instance routing-instance-name;
        static-interface interface-name;
      }
    }
  }
}
```

Hierarchy Level [edit protocols pppoe]

Release Information Statement introduced in Junos OS Release 10.0.
dynamic-profile, **routing-instance**, **max-sessions**, and **static-interface** options introduced in Junos OS Release 10.2.

Description Create and configure a PPPoE service name table. Specify the action taken for each service and remote access concentrator on receipt of a PPPoE Active Discovery Initiation (PADI) packet. You can also specify the dynamic profile and routing instance that the router uses to instantiate a dynamic PPPoE interface, and the maximum number of active PPPoE sessions that the router can establish with the specified service. A maximum of 32 PPPoE service name tables is supported per router.

Options ***table-name***—Name of the PPPoE service name table, a string of up to 32 alphanumeric characters.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring PPPoE Service Name Tables*
- *Creating a Service Name Table*

service-set

Syntax

service-set *service-set-name*;

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* family inet service (input | output) service-set *service-set-name*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet service (input | output) service-set *service-set-name*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Define one or more service sets to be applied to an interface. If you define multiple service sets, the Junos OS evaluates the filters in the order in which they appear in the configuration.

Options

service-set-name—Identifies the service set.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

 Copyright © 2018, Juniper Networks, Inc.

services (Priority Level)

Syntax services *priority-level*;

Hierarchy Level [edit interfaces *interface-name* services-options sysloghost *hostname*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify system logging priority level.

Options *priority-level*—Assigns a priority level to the facility. Valid entries are as follows:

- **alert**—Conditions that should be corrected immediately.
- **any**—Matches any level.
- **emergency**—Panic conditions.
- **critical**—Critical conditions.
- **error**—Error conditions.
- **info**—Informational messages.
- **notice**—Conditions that require special handling.
- **warning**—Warning messages.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2018, Juniper Networks, Inc.

services-options

Syntax

```
services-options {
  inactivity-timeout seconds;
  open-timeout seconds;
    session-limit {
      maximum number;
      rate new-sessions-per-second;
    }
  syslog {
    host hostname {
      facility-override facility-name;
      log-prefix prefix-number;
      services priority-level;
    }
  }
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the service options to be applied on an interface.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

 Copyright © 2018, Juniper Networks, Inc.

shaping

Syntax

shaping {
  (cbr *rate* | rtvbr peak *rate* sustained *rate* burst *length* | vbr peak *rate* sustained *rate* burst
    length);
  queue-length *number*;
}

Hierarchy Level

[edit interfaces *interface-name* atm-options vpi *vpi-identifier*],
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit interfaces *interface-name* unit *logical-unit-number* address *address* family *family*
  multipoint-destination *address*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*
  address *address* family *family* multipoint-destination *address*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

For ATM encapsulation only, define the traffic-shaping profile.

For Circuit Emulation PICs, specify traffic shaping in the ingress and egress directions.

For ATM2 IQ interfaces, changing or deleting VP tunnel traffic shaping causes all logical interfaces on a VP to be deleted and then re-added.

VP tunnels are not supported on multipoint interfaces.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Defining Virtual Path Tunnels*
- *Defining the ATM Traffic-Shaping Profile Overview*
- *Configuring ATM QoS or Shaping*
- *Applying Scheduler Maps to Logical ATM Interfaces*

shdsl-options

Syntax

```
shdsl-options {
  annex (annex-a | annex-b);
  line-rate line-rate;
  loopback (local | remote | payload);
  snr-margin {
    snext margin;
  }
}
```

Hierarchy Level

[edit interfaces *interface-name*],
[edit logical-systems *logical-system-name* interfaces *interface-name*

Release Information

Statement introduced in Junos OS Release 7.4.

Description

For J Series Services Routers only, configure symmetric DSL (SHDSL) options.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2018, Juniper Networks, Inc.

short-name-format

Syntax short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name*]

Release Information Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

Description Specify the name format of the maintenance association name.

Options **character-string**—The name is an ASCII character string.

vlan—The primary VLAN identifier.

2octet—A number in the range 0 through 65,535.

rfc-2685-vpn-id—A VPN identifier that complies with RFC 2685.

Default: character-string

NOTE: The PTX Series Packet Transport Routers support the **vlan** and **2octet** options only.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Creating a Maintenance Association*

Copyright © 2018, Juniper Networks, Inc.

short-sequence

Syntax short-sequence;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For multilink interfaces only, set the length of the packet sequence identification number to 12 bits.

Default If you omit this statement from the configuration, the length is set to 24 bits.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Junos OS Services Interfaces Library for Routing Devices*

 Copyright © 2018, Juniper Networks, Inc.

sla-iterator-profile

Syntax

```
sla-iterator-profile profile-name {
  data-tlv-size size;
  iteration-count count-value;
  priority priority-value;
}
```

Hierarchy Level

[edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id* remote-mep *remote-mep-id*]

Release Information

Statement introduced in Junos OS Release 11.1.

Description

Configure a remote MEP with an iterator profile and specify the options.

Options

profile-name—Name of the iterator profile configured for a remote MEP. For more information about configuring a remote MEP with an iterator profile, see *Configuring a Remote MEP with an Iterator Profile*.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- *Configuring an Iterator Profile*
- *Configuring a Remote MEP with an Iterator Profile*
- *Verifying the Configuration of an Iterator Profile*
- *Managing Iterator Statistics*
- sla-iterator-profiles on page 970

Copyright © 2018, Juniper Networks, Inc.

sla-iterator-profiles

Syntax

```
sla-iterator-profiles {
  profile-name {
    avg-fd-twoway-threshold;
      avg-ifdv-twoway-threshold;
      avg-flr-forward-threshold;
      avg-flr-backward-threshold;
      calculation-weight {
      delay delay-weight;
      delay-variation delay-variation-weight;
    }
    cycle-time milliseconds;
    flap-trap-monitor seconds
    iteration-period iteration-period-value;
    measurement-type (loss | statistical-frame-loss | two-way-delay);
  }
}
```

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management performance-monitoring]

Release Information Statement introduced in Junos OS Release 11.1.

Description Configure an iterator application and specify the iterator profile options.

Options *profile-name*—Name of the iterator profile. For more information about configuring the iterator profile, see *Configuring an Iterator Profile*.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- *Configuring an Iterator Profile*
- *Configuring a Remote MEP with an Iterator Profile*
- *Verifying the Configuration of an Iterator Profile*
- *Managing Iterator Statistics*

 Copyright © 2018, Juniper Networks, Inc.

snr-margin

Syntax

```
snr-margin {
  snext margin;
}
```

Hierarchy Level

[edit interfaces *interface-name* shdsl-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* shdsl-options]

Release Information

Statement introduced in Junos OS Release 7.4.

Description

For J Series Services Routers only, configure the SHDSL signal-to-noise ratio (SNR) margin. The SNR margin is the difference between the desired SNR and the actual SNR. Configuring the SNR creates a more stable SHDSL connection by making the line train at a SNR margin higher than the threshold. If any external noise below the threshold is applied to the line, the line remains stable.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Junos OS Interfaces and Routing Configuration Guide*

Copyright © 2018, Juniper Networks, Inc.

snext

Syntax snext *margin*;

Hierarchy Level [edit interfaces *interface-name* shdsl-options snr-margin],
[edit logical-systems *logical-system-name* interfaces *interface-name* shdsl-options snr-margin]

Release Information Statement introduced in Junos OS Release 7.4.

Description For J Series Services Routers only, configure self-near-end crosstalk (SNEXT) signal-to-noise ratio (SNR) margin for a SHDSL line. When configured, the line trains at higher than SNEXT threshold. The SNR margin is the difference between the desired SNR and the actual SNR.

Options *margin*—Desired SNEXT margin. Possible values are disabled or a margin between –10dB and 10 dB.

Default: disabled

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Junos OS Interfaces and Routing Configuration Guide*

Copyright © 2018, Juniper Networks, Inc.

sonet-options

Syntax

```
sonet-options {
  aps {
    advertise-interval milliseconds;
    annex-b
    authentication-key key;
    (break-before-make | no-break-before-make);
    fast-aps-switch;
    force;
    hold-time milliseconds;
    lockout;
    neighbor address;
    paired-group group-name;
    protect-circuit group-name;
    request;
    revert-time seconds;
    switching-mode (bidirectional | unidirectional);
    working-circuit group-name;
  }
  bytes {
    c2 value;
    e1-quiet value;
    f1 value;
    f2 value;
    s1 value;
    z3 value;
    z4 value;
  }
  fcs (16 | 32);
  loopback (local | remote);
  mpls {
    pop-all-labels {
      required-depth number;
    }
  }
  path-trace trace-string;
  (payload-scrambler | no-payload-scrambler);
  rfc-2615;
  trigger {
    defect ignore;
    defect hold-time up milliseconds down milliseconds;
  }
}
vtmapping (itu-t | klm);
(z0-increment | no-z0-increment);
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure SONET/SDH-specific interface properties.

Copyright © 2018, Juniper Networks, Inc.

On SONET/SDH OC48 interfaces that you configure for channelized (multiplexed) mode (by including the **no-concatenate** statement at the **[edit chassis fpc *slot-number* pic *pic-number*]** hierarchy level), the **bytes e1-quiet** and **bytes f1** options have no effect. The **bytes f2**, **bytes z3**, **bytes z4**, and **path-trace** options work correctly on channel 0 and work in the transmit direction only on channels 1, 2, and 3.

On a channelized OC12 interface, the **bytes e1-quiet, bytes f1, bytes f2, bytes z3,** and **bytes z4** options are not supported. The **fcs** and **payload-scrambler** statements are also not supported; you must configure these for each DS3 channel using the **t3-options fcs** and **t3-options payload-scrambler** statements. The **aps** and **loopback** statements are supported only on channel 0 and are ignored if included in the configurations for channels 1 through 11. You can configure loopbacks for each DS3 channel with the **t3-options loopback** statement. The **path-trace** statement can be included in the configuration for each DS3 channel, thereby configuring a unique path trace for each channel.

To configure loopback on channelized IQ and IQE PICs, SONET/SDH level, use the **loopback** statement **local** and **remote** options at the controller interface (coc48, cstm16, coc12, cstm4, coc3, and cstm1). It is ignored for path-level interfaces **so-*fpc/pic/port*** or **so-*fpc/pic/port*:*channel***.

If you are running Intermediate System-to-Intermediate System (IS-IS) over SONET/SDH interfaces, use PPP if you are running Cisco IOS Release 12.0 or later. If you need to run HDLC, configure an ISO family MTU of 4469 on the router.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring SONET/SDH Parameters on ATM Interfaces*
- *Channelized OC12/STM4 IQ and IQE Interfaces Overview*
- *Channelized STM1 Interfaces Overview*
- *SONET/SDH Interfaces Overview*
- *no-concatenate*

 Copyright © 2018, Juniper Networks, Inc.

source

Syntax source *source-address*;

Hierarchy Level (EX, NFX, OCX1100 and QFX Series) [edit interfaces *interface-name* unit *logical-unit-number* tunnel]

Hierarchy Level (M-series and T-series) [edit interfaces *interface-name* unit *logical-unit-number* tunnel *address*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* tunnel *address*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Specify the source address of the tunnel.

Default If you do not specify a source address, the tunnel uses the unit's primary address as the source address of the tunnel.

Options *source-address*—Address of the local side of the tunnel. This is the address that is placed in the outer IP header's source field.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Tunnel Services Overview*
- multicast-only on page 791
- primary (Address on Interface) on page 901
- *Junos OS Services Interfaces Library for Routing Devices*

source-address-filter

Syntax

source-address-filter {
 mac-address;
}

Hierarchy Level

[edit interfaces *interface-name* aggregated-ether-options],
[edit interfaces *interface-name* fastether-options],
[edit interfaces *interface-name* gigether-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Packet Transport Routers.

Description

For aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, Gigabit Ethernet IQ interfaces, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), specify the MAC addresses from which the interface can receive packets. For this statement to have any effect, you must include the **source-filtering** statement in the configuration to enable source address filtering.

Options

mac-address—MAC address filter. You can specify the MAC address as ***nn*:*nn*:*nn*:*nn*:*nn*:*nn*** or ***nnnn*.*nnnn*.*nnnn***, where ***n*** is a decimal digit. To specify more than one address, include multiple ***mac-address*** options in the **source-address-filter** statement.

If you enable the VRRP on a Fast Ethernet or Gigabit Ethernet interface, as described in *VRRP and VRRP for IPv6 Overview*, and if you enable MAC source address filtering on the interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the **source-address-filter** statement. MAC addresses ranging from **00:00:5e:00:01:00** through **00:00:5e:00:01:ff** are reserved for VRRP, as defined in RFC 3768, *Virtual Router Redundancy Protocol*. When you configure the VRRP group, the group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.

On untagged Gigabit Ethernet interfaces, you should not configure the **source-address-filter** statement and the **accept-source-mac** statement simultaneously. On tagged Gigabit Ethernet interfaces, you should not configure the **source-address-filter** statement and the **accept-source-mac** statement with an identical MAC address specified in both filters.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring MAC Address Filtering for Ethernet Interfaces*
- *Configuring MAC Address Filtering on PTX Series Packet Transport Routers*
- source-filtering on page 826

 Copyright © 2018, Juniper Networks, Inc.

source-class-usage

Syntax

source-class-usage {
  direction;
}

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* family inet accounting],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet accounting],
[edit routing-instances *routing-instance-name* vrf-table-label]

Release Information

Statement introduced before Junos OS Release 7.4.
Support for the **vrf-table-label** statement added in Junos OS Release 9.3.

Description

Enable packet counters on an interface that count packets that arrive from specific prefixes on the provider core router and are destined for specific prefixes on the customer edge router.

Options

direction can be one of the following:

input—Configure at least one expected ingress point.

output—Configure at least one expected egress point.

input output—On a single interface, configure at least one expected ingress point and one expect egress point.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Enabling Source Class and Destination Class Usage on page 262
- **accounting on page 413**
- **destination-class-usage on page 530**
- *Junos OS Services Interfaces Library for Routing Devices*
- *vrf-table-label*

Copyright © 2018, Juniper Networks, Inc.

source-filtering

Syntax (source-filtering | no-source-filtering);

Hierarchy Level [edit interfaces *interface-name* aggregated-ether-options],
[edit interfaces *interface-name* fastether-options],
[edit interfaces *interface-name* gigether-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Packet Transport Routers.

Description For aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, and Gigabit Ethernet IQ interfaces only, enable the filtering of MAC source addresses, which blocks all incoming packets to that interface. To allow the interface to receive packets from specific MAC addresses, include the **source-address-filter** statement.

If the remote Ethernet card is changed, the interface is no longer able to receive packets from the new card because it has a different MAC address.

Default Source address filtering is disabled.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring MAC Address Filtering for Ethernet Interfaces*
- *Configuring MAC Address Filtering on PTX Series Packet Transport Routers*
- accept-source-mac on page 409
- source-address-filter on page 976

 Copyright © 2018, Juniper Networks, Inc.

speed (Ethernet)

List of Syntax

Syntax (EX Series)

speed (auto-negotiation | *speed*) ;

Syntax (EX2300 and EX4300)

speed *speed*;

Syntax (EX Series, ACX Series, MX Series)

speed (10m |10g | 100m | 1g | 2.5g |5g |auto | auto-10m-100m);

Syntax (QFX Series, OCX1100, EX4600)

speed (10g | 1g | 100m)

Hierarchy Level (EX Series)

[edit interfaces *interface-name* ether-options]

Hierarchy Level (EX2300 and EX4300)

[edit interfaces *interface-name*]

Hierarchy Level (ACX Series, EX Series, MX Series)

[edit interfaces *interface-name*],
[edit interfaces ge-*pim*/0/0 switch-options switch-port *port-number*]

Hierarchy Level (QFX Series, EX4600, OCX Series)

[edit interfaces *interface-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
Speed option 2.5Gbps introduced in Junos OS Release 18.1R2 for EX2300 switch.
Speed option 10Gbps and 5Gbps introduced in Junos OS Release 18.2R1 for EX4300 switch.

Description

Configure the interface speed. This statement applies to the management Ethernet interface (**fxp0** or **em0**), Fast Ethernet 12-port and 48-port PICs, the built-in Fast Ethernet port on the FIC (M7i router), Combo Line Rate DPCs and Tri-Rate Ethernet Copper interfaces on MX Series routers, and Gigabit Ethernet interfaces on EX Series switches.

When you configure the Tri-Rate Ethernet copper interface to operate at 1 Gbps, autonegotiation must be enabled. When you configure 100BASE-FX SFP, you must set the port speed at 100 Mbps.

Copyright © 2018, Juniper Networks, Inc.

NOTE: On MX Series routers with Tri-rate Enhanced DPC (DPCE-R-40GE-TX), when you configure the interface speed using the **auto-10m-100m** option, the speed is negotiated to the highest value possible (100 Mbps), if the same value is configured on both sides of the link. However, when you view the interface speed of the DPC, using the **show interfaces** command, the value of the speed is not accurately displayed. For instance, if you configure the speed of the Tri-rate enhanced DPC, as 100Mbps on both sides of the link, the interface speed of the DPC is negotiated to 100 Mbps. However, the interface speed of the DPC displays 1 bps. This is an issue with the **show interfaces** command only. The actual interface speed is 100 Mbps.

On 10-Gigabit Ethernet SFP interfaces, autonegotiation is enabled by default and auto-detects the speed to be either 1 Gbps or 10 Gbps. On QFX5100-48S, QFX5100-96S, and QFX5100-24Q devices using 10-Gigabit Ethernet SFP interfaces, the speed is set to 10 Gbps by default and cannot be configured to operate in a different speed. On QFX5100-48S and QFX5100-96S devices using 1-Gigabit Ethernet SFP interfaces, the speed is set to 1 Gbps by default and cannot be configured to operate in a different speed.

NOTE: In Junos OS Release 14.1X53-D35 on QFX5100-48T-6Q devices using 10-Gigabit Ethernet Copper interfaces, autonegotiation is disabled by default on the copper ports, and the interfaces operate at a speed of 100M. You can, however, enable auto-negotiation by issuing the **set interface *name* ether-options auto-negotiation** command on the interface for which you want to change the interface speed. With autonegotiation enabled, the interface auto-detects the speed in which to operate.

NOTE: Only 10 Gbps and 40 Gbps interfaces are supported on OCX Series switches.

NOTE: When displaying interface information with **show interfaces** commands, you might see speed values for 1 Gbps interfaces displayed as **1000mbps**.

 Copyright © 2018, Juniper Networks, Inc.

(For EX2300 only) Starting in Junos OS Release 18.1R2, the multi-rate speed is supported on EX2300-48MP and EX2300-24MP switches. The speed configuration statement is supported on both multi-rate gigabit ethernet interface (mge) and gigabit ethernet (ge) interface. The mge interface is a rate-selectable (multirate) Gigabit Ethernet interface that can support speeds of 10 Gbps, 5 Gbps, and 2.5 Gbps over CAT5e/CAT6/CAT6a cables. In the EX2300, the mge interface supports 100 Mbps, 1 Gbps, and 2.5 Gbps speeds, which can be configured by using the speed configuration statement. Note that 10Mbps speed is supported only on **ge** interfaces of EX2300 switch.

Default (EX Series) If the auto-negotiation statement at the **[edit interfaces *interface-name* ether-options]** hierarchy level is enabled, the auto-negotiation option is enabled by default.

Options You can specify the speed as either **10m** (10 Mbps), **100m** (100 Mbps), and on MX Series routers, **1g** (1 Gbps). You can also specify the **auto** option on MX Series routers.

For Gigabit Ethernet interfaces on EX Series switches, you can specify one of the following options:

Table 53: Options for speed

Platforms	Speed Supported	Auto-negotiation
EX Series Switches	**100m**—100 Mbps **10m**—10 Mbps **1g**—1 Gbps	auto-negotiation—Automatically negotiate the speed based on the speed of the other end of the link. This option is available only when the auto-negotiation statement at the **[edit interfaces *interface-name* ether-options]** hierarchy level is enabled.
ACX, MX Series	**100m**—100 Mbps **10m**—10 Mbps **1g**—1 Gbps	**auto**—Automatically negotiate the speed (10 Mbps, 100 Mbps, or 1 Gbps) based on the speed of the other end of the link. **auto-10m-100m**—Automatically negotiate the speed (10 Mbps or 100 Mbps) based on the speed of the other end of the link.
EX4600, QFX Series, QFabric, OCX100, QFX Series	**10g**—10 Gbps **1g**—1 Gbps **100m**—100 Mbps	auto-negotiation—Automatically negotiate the speed based on the speed of the other end of the link. This option is available only when the auto-negotiation statement at the **[edit interfaces *interface-name* ether-options]** hierarchy level is enabled.
EX2300	**10m**—10 Mbps (supported on EX series switches and only on **ge** interfaces of EX2300 switch) **100m**—100 Mbps **1g**—1 Gbps **2.5g**—2.5 Gbps (supported only on **mge** interfaces of E2300 switch) **10g**—10 Gbps (supported only on **mge** interfaces for EX4300 switches) **5g**—5 Gbps (supported only on **mge** interfaces for EX4300 switches)	*speed*—Specify the interface speed. If the auto-negotiation statement at the **[edit interfaces *interface-name* ether-options]** hierarchy level is disabled, you must specify a specific value. This value sets the speed that is used on the link. If the **auto-negotiation** statement is enabled, you might want to configure a specific speed value to advertise the desired speed to the remote end.

Copyright © 2018, Juniper Networks, Inc.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release History Table

Release	Description
18.2R1	Starting in Junos OS Release 18.1R2, the multi-rate speed is supported on EX2300-48MP and EX2300-24MP switches.

Related Documentation

- Configuring the Interface Speed on page 116
- Configuring the Interface Speed on Ethernet Interfaces on page 116
- *Configuring Gigabit Ethernet Autonegotiation*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure) for EX Series Switches with ELS support*
- *auto-negotiation*
- *Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX and QFX Series Switches*
- *Junos OS Network Interfaces Library for Routing Devices*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit Ethernet Interfaces (J-Web Procedure)*
- *Junos OS Ethernet Interfaces Configuration Guide*

Copyright © 2018, Juniper Networks, Inc.

speed (MX Series DPC)

Syntax speed (auto | 1Gbps | 100Mbps | 10Mbps);

Hierarchy Level [edit interfaces ge-*fpc/pic/port*]

Release Information Statement introduced in Junos OS Release 9.5.

Description On MX Series routers with Combo Line Rate DPCs and Tri-Rate Copper SFPs you can set auto negotiation of speed. To specify the auto negotiation speed, use the **speed (auto | 1Gbps | 100Mbps | 10Mbps)** statement under the **[edit interface ge-/fpc/pic/port]** hierarchy level. The **auto** option will attempt to automatically match the rate of the connected interface. To set port speed negotiation to a specific rate, set the port speed to **1Gbps**, **100Mbps**, or **10Mbps**.

NOTE: If the negotiated speed and the interface speed do not match, the link will not be brought up. Half duplex mode is not supported.

Options You can specify the speed as either **auto** (autonegotiate), **10Mbps** (10 Mbps), **100Mbps** (100 Mbps), or **1Gbps** (1 Gbps).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Gigabit Ethernet Autonegotiation*
- no-auto-mdix on page 809

speed (SONET/SDH)

Syntax speed (oc3 | oc12 | oc48);

Hierarchy Level [edit interfaces so-*fpc/pic/port*],
[edit interfaces so-*fpc/pic/port*:*channel*]

Release Information Statement introduced in Junos OS Release 8.3.

Description Configure the interface speed. This statement applies to SONET/SDH interfaces on next-generation SONET/SDH Type 1 and Type 2 PICs with SFP. Available speeds depend on whether the PIC is in concatenated mode or nonconcatenated mode. Include the channel in the interface name when configuring nonconcatenated interfaces.

Options **oc3 | oc12 | oc48**—Speed when the PIC is in concatenated mode. For example, you can configure each port of a 4-port OC12 PIC to have a speed of **oc3**.

You can configure port 0 of a 4-port OC12 PIC to have a speed of **oc12**.

oc3 | oc12—Speed when the PIC is in nonconcatenated mode.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring SONET/SDH Interface Speed on page 120

spid1

Syntax spid1 *spid1-string*;

Hierarchy Level [edit interfaces br-*pim*/0/*port* isdn-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the Service Profile Identifier (SPID).

Options ***spid1-string***—Numeric SPID.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Junos OS Interfaces and Routing Configuration Guide*

 Copyright © 2018, Juniper Networks, Inc.

spid2

Syntax spid2 *spid2-string*;

Hierarchy Level [edit interfaces br-*pim*/0/*port* isdn-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure an additional SPID.

Options ***spid2-string***—Numeric SPID.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

Copyright © 2018, Juniper Networks, Inc.

stacked-vlan-ranges

Syntax

```
stacked-vlan-ranges {
  access-profile profile-name;
  authentication {
    packet-types [packet-types];
    password password-string;
    username-include {
      circuit-type;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-name;
      mac-address;
      option-18
      option-37
      option-82;
      radius-realm radius-realm-string;
      user-prefix user-prefix-string;
    }
  }
  dynamic-profile profile-name {
    accept (any | dhcp-v4 | inet);
    access-profile vlan-dynamic-profile-name;
    ranges (any | low-tag–high-tag),(any | low-tag–high-tag);
  }
  override;
}
```

Hierarchy Level [edit interfaces *interface-name* auto-configure]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure multiple VLANs. Each VLAN is assigned a VLAN ID number from the range.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation

- *Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs*
- *Configuring Interfaces to Support Both Single and Stacked VLANs*

 Copyright © 2018, Juniper Networks, Inc.

stacked-vlan-tagging

Syntax stacked-vlan-tagging;

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description For Gigabit Ethernet IQ interfaces, Gigabit Ethernet, 10-Gigabit Ethernet LAN/WAN PIC, and 100-Gigabit Ethernet Type 5 PIC with CFP, enable stacked VLAN tagging for all logical interfaces on the physical interface.

For pseudowire subscriber interfaces, enable stacked VLAN tagging for logical interfaces on the pseudowire service.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview*

Copyright © 2018, Juniper Networks, Inc.

start-end-flag

Syntax start-end-flag (filler | shared);

Hierarchy Level [edit interfaces e1-*fpc/pic/port*],
[edit interfaces t1-*fpc/pic/port*],
[edit interfaces *interface-name* ds0-options],
[edit interfaces *interface-name* e1-options],
[edit interfaces *interface-name* e3-options],
[edit interfaces *interface-name* t1-options],
[edit interfaces *interface-name* t3-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers.

Description For DS0, E1, E3, T1, and T3 interfaces, configure the interface to share the transmission of start and end flags.

NOTE: When configuring E1 or T1 interfaces on the 10-port Channelized E1/T1 IQE PIC, the **start-end-flag** statement must be included at the **[edit interfaces e1-*fpc/pic/port*]** or **[edit interfaces t1-*fpc/pic/port*]** hierarchy level as appropriate.

Options **filler**—Wait two idle cycles between the start and end flags.

shared—Share the transmission of the start and end flags. This is the default.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring E1 Start and End Flags*
- *Configuring the E3 Start and End Flags*
- *Configuring T1 Start and End Flags*
- *Configuring T3 Start and End Flags*

 Copyright © 2018, Juniper Networks, Inc.

static-interface

Syntax static-interface *interface-name*;

Hierarchy Level [edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Reserve the specified static PPPoE interface for use only by the PPPoE client with matching agent circuit identifier (ACI) and agent remote identifier (ARI) information. You can specify only one static interface per ACI/ARI pair configured for a named service entry, **empty** service entry, or **any** service entry in the PPPoE service name table.

The static interface associated with an ACI/ARI pair takes precedence over the general pool of static interfaces associated with the PPPoE underlying interface.

If you include the **static-interface** statement in the configuration, you cannot also include either the **dynamic-profile** statement or the **routing-instance** statement. The **dynamic-profile**, **routing-instance**, and **static-interface** statements are mutually exclusive for ACI/ARI pair configurations.

Options *interface-name*—Name of the static PPPoE interface reserved for use by the PPPoE client with matching ACI/ARI information. Specify the interface in the format **pp0.*logical***, where ***logical*** is a logical unit number from 0 through 16385 for static interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring PPPoE Service Name Tables*
- *Reserving a Static PPPoE Interface for Exclusive Use by a PPPoE Client*

Copyright © 2018, Juniper Networks, Inc.

static-tei-val

Syntax static-tei-val *value*;

Hierarchy Level [edit interfaces br-*pim*/0/*port* isdn-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For J Series Services Routers only. Statically configure the Terminal Endpoint Identifier (TEI) value. The TEI value represents any ISDN-capable device attached to an ISDN network that is the terminal endpoint. TEIs are used to distinguish between several different devices using the same ISDN links.

Options ***value***—Value between 0 through 63.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Junos OS Interfaces and Routing Configuration Guide*

 Copyright © 2018, Juniper Networks, Inc.

supplicant

Syntax supplicant *single*;

Hierarchy Level [edit protocols dot1x authenticator interface *interface-id*]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the supplicant mode. Only single mode is supported.

This option will authenticate only the first client that connects to a port. All other clients that connect later (802.1x compliant or non-compliant) will be allowed free access on that port without any further authentication. If the first authenticated client logs out, all other users are locked out until a client authenticates again.

Options **single**—Sets single mode.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *IEEE 802.1x Port-Based Network Access Control Overview*
- authenticator on page 448
- dot1x on page 543
- interface (IEEE 802.1x) on page 682

Copyright © 2018, Juniper Networks, Inc.

supplicant-timeout

Syntax supplicant-timeout *seconds*;

Hierarchy Level [edit protocols dot1x authenticator interface *interface-id*]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the number of seconds the port waits for a response when relaying a request from the authentication server to the client before resending the request.

Options *seconds*—Specify the number of seconds the port waits for the supplicant timeout.

Range: 1 through 60 seconds

Default: 30 seconds

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *IEEE 802.1x Port-Based Network Access Control Overview*
- authenticator on page 448
- dot1x on page 543
- interface (IEEE 802.1x) on page 682

 Copyright © 2018, Juniper Networks, Inc.

swap

Syntax swap;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.

Description Specify the VLAN rewrite operation to replace a VLAN tag. The outer VLAN tag of the frame is overwritten with the user-specified VLAN tag information.

On MX Series routers, you can enter this statement on Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, aggregated Ethernet using Gigabit Ethernet IQ interfaces, and 100-Gigabit Ethernet Type 5 PIC with CFP. On EX Series switches, you can enter this statement on Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Rewriting the VLAN Tag on Tagged Frames*
- *Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support (CLI Procedure)*

swap-push

Syntax swap-push;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify the VLAN rewrite operation to replace the outer VLAN tag of the frame with a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame.

You can use this statement on Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Rewriting a VLAN Tag and Adding a New Tag*

 Copyright © 2018, Juniper Networks, Inc.

swap-swap

Syntax

swap-swap;

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information

Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Specify the VLAN rewrite operation to replace both the inner and the outer VLAN tags of the frame with a user-specified VLAN tag value.

You can use this statement on Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and for 100-Gigabit Ethernet Type 5 PIC with CFP.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Rewriting the Inner and Outer VLAN Tags*

switch-options

Syntax

```
switch-options {
  switch-port port-number {
    (auto-negotiation | no-auto-negotiation);
    speed (10m | 100m | 1g);
    link-mode (full-duplex | half-duplex);
  }
}
```

Hierarchy Level

[edit interfaces ge-*pim*/0/0]

Release Information

Statement introduced in Junos OS Release 8.4.

Description

Configuration of the physical port characteristics is done under the single physical interface.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

switch-port

Syntax

switch-port *port-number* {
(auto-negotiation | no-auto-negotiation);
speed (10m | 100m | 1g);
link-mode (full-duplex | half-duplex);
}

Hierarchy Level

[edit interfaces ge-*pim*/0/0 switch-options]

Release Information

Statement introduced in Junos OS Release 8.4.

Description

Configuration of the physical port characteristics, done under the single physical interface.

Default

Autonegotiation is enabled by default. If the link speed and duplex are also configured, the interfaces use the values configured as the desired values in the negotiation.

Options

port-number—Ports are numbered 0 through 5 on the 6-port Gigabit Ethernet uPIM, 0 through 7 on the 8-port Gigabit Ethernet uPIM, and 0 through 15 on the 16-port Gigabit Ethernet uPIM.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2018, Juniper Networks, Inc.

switch-type

Syntax switch-type (att5e | etsi | ni1 | ntdms-100)

Hierarchy Level [edit interfaces br-*pim*/0/*port* isdn-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For J Series Services Routers only. Configure the ISDN variant supported.

Options **att5e**—AT&T switch variant.

etsi—European Telecommunications Standards Institute switch variant.

ni1—National ISDN 1 switch variant.

ntdms-100—Northern Telecom DMS-100.

ntt—NTT Group switch for Japan.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Junos OS Interfaces and Routing Configuration Guide*

Copyright © 2018, Juniper Networks, Inc.

switching-mode

Syntax switching-mode (bidirectional | unidirectional);

Hierarchy Level [edit interfaces *interface-name* sonet-options aps]

Release Information Statement introduced before Junos OS Release 7.4.

Description For unchannelized OC3, OC12, and OC48 SONET/SDH interfaces on T Series routers only, configure the interface to interoperate with SONET/SDH line-terminating equipment (LTE) that is provisioned for unidirectional linear APS in 1+1 architecture.

Default If the **switching-mode** statement is not configured, the mode is bidirectional, and the interface does not interoperate with a unidirectional SONET/SDH LTE.

Options **bidirectional**—Support bidirectional mode only.

unidirectional—Interoperate with a SONET/SDH LTE provisioned for unidirectional mode.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Switching Mode*

 Copyright © 2018, Juniper Networks, Inc.

symbol-period

Syntax

symbol-period *count*;

Hierarchy Level

[edit protocols oam ethernet link-fault-management action-profile event, link-event-rate],
[edit protocols oam link-fault-management interface *interface-name* event-thresholds]

Release Information

Statement introduced in Junos OS Release 8.4.

Description

Configure the threshold for sending symbol period events or taking the action specified in the action profile.

A symbol error is any symbol code error on the underlying physical layer. The symbol period threshold is reached when the number of symbol errors reaches the configured value within the period window. The default period window is the number of symbols that can be transmitted on the underlying physical layer in 1 second. The window is not configurable.

Options

count—Threshold count for symbol period events.

Range: 0 through 100

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Threshold Values for Local Fault Events on an Interface*
- *Configuring Threshold Values for Fault Events in an Action Profile*

Copyright © 2018, Juniper Networks, Inc.

syslog (Interfaces)

Syntax

```
syslog {
  host hostname {
    facility-override facility-name;
    log-prefix prefix-number;
    services priority-level;
  }
}
```

Hierarchy Level [edit interfaces *interface-name* services-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For adaptive services interfaces, configure generation of system log messages for the service set. System log information is passed to the kernel for logging in the **/var/log** directory. Any values configured in the service set definition override these values.

Options The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

syslog (Monitoring)

Syntax (syslog | no-syslog);

Hierarchy Level [edit interfaces mo-*fpc/pic/port* multiservice-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description System logging is enabled by default. The system log information of the Monitoring Services PIC is passed to the kernel for logging in the **/var/log** directory.

- **syslog**—Enable PIC system logging.
- **no-syslog**—Disable PIC system logging.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Multiservice Physical Interface Properties on page 170
- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2018, Juniper Networks, Inc.

syslog (OAM Action)

Syntax syslog;

Hierarchy Level [edit protocols oam ethernet link-fault-management action-profile action]

Release Information Statement introduced in Junos OS Release 8.5 for T, M, MX and ACX Series routers, SRX Series firewalls and EX Series switches.
Statement introduced in Junos OS Release 9.4 for EX Series switches and NFX Series devices.

Description Generate a syslog message for the Ethernet Operation, Administration, and Management (OAM) event.

Generate a system log message for the Ethernet Operation, Administration, and Maintenance (OAM) link fault management (LFM) event.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Specifying the Actions to Be Taken for Link-Fault Management Events*
- *Configuring Ethernet OAM Link Fault Management (CLI Procedure)*

Copyright © 2018, Juniper Networks, Inc.

system-priority

Syntax system-priority *priority*;

Hierarchy Level [edit interfaces ae*X* aggregated-ether-options lacp]

Release Information Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 11.4 for EX Series switches.
Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.

Description Define LACP system priority at the aggregated Ethernet interface level. This system priority value takes precedence over a system priority value configured at the global **[edit chassis]** hierarchy level.

The device with the lower system priority value determines which links between LACP partner devices are active and which are in standby for each LACP group. The device on the controlling end of the link uses port priorities to determine which ports are bundled into the aggregated bundle and which ports are put in standby mode. Port priorities on the other device (the noncontrolling end of the link) are ignored. In priority comparisons, numerically lower values have higher priority. Therefore, the system with the numerically lower value (higher priority value) for LACP system priority becomes the controlling system. If both devices have the same LACP system priority (for example, they are both configured with the default setting of 127), the device MAC address determines which switch is in control.

Options ***priority***—Priority for the aggregated Ethernet system. A smaller value indicates a higher priority.

Range: 0 through 65535

Default: 127

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2018, Juniper Networks, Inc.

t1-options

Syntax

```
t1-options {
  bert-algorithm algorithm;
  bert-error-rate rate;
  bert-period seconds;
  buildout value;
  byte-encoding (nx56 | nx64);
  crc-major-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5);
  crc-minor-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5 | 5e-6 | 1e-6);
  fcs (16 | 32);
  framing (esf | sf);
  idle-cycle-flag (flags | ones);
  invert-data;
  line-encoding (ami | b8zs);
  loopback (local | payload | remote);
  remote-loopback-respond;
  start-end-flag (filler | shared);
  timeslots time-slot-range;
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers.

Description Configure T1-specific physical interface properties.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *T1 Interfaces Overview*

t310

Syntax t310-value *seconds*;

Hierarchy Level [edit interfaces br-*pim*/0/*port* isdn-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For ISDN interfaces, configure the Q.931-specific timer for T310, in seconds. The Q.931 protocol is involved in the setup and termination of connections.

Options ***seconds***—Timer value, in seconds.

Range: 1 through 65,536 seconds

Default: 10 seconds

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Junos OS Interfaces and Routing Configuration Guide*

 Copyright © 2018, Juniper Networks, Inc.

t391

Syntax t391 *seconds*;

Hierarchy Level [edit interfaces *interface-name* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services and voices interfaces only, set Frame Relay link integrity polling interval.

Options *seconds*—Link integrity polling interval.

Range: 5 through 30 seconds

Default: 10 seconds

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- n391 on page 796
- n392 on page 797
- n393 on page 798
- t392 on page 1006

- *Junos OS Services Interfaces Library for Routing Devices*

t392

Syntax t392 *seconds*;

Hierarchy Level [edit interfaces *interface-name* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services and voices interfaces only, set Frame Relay polling verification interval.

Options *seconds*—Polling verification interval.

Range: 5 through 30 seconds

Default: 15 seconds

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- n391 on page 796
- n392 on page 797
- n393 on page 798
- t391 on page 1005
- timeslots on page 1016

- *Junos OS Services Interfaces Library for Routing Devices*

 Copyright © 2018, Juniper Networks, Inc.

t3-options

Syntax

```
t3-options {
  atm-encapsulation (direct | plcp);
  bert-algorithm algorithm;
  bert-error-rate rate;
  bert-period seconds;
  (cbit-parity | no-cbit-parity);
  compatibility-mode (digital-link | kentrox | larscom) <subrate value>;
  fcs (16 | 32);
  (feac-loop-respond | no-feac-loop-respond);
  idle-cycle-flag value;
  (long-buildout | no-long-buildout);
  (loop-timing | no-loop-timing);
  loopback (local | payload | remote);
  start-end-flag value;
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure T3-specific physical interface properties, including the properties of DS3 channels on a channelized OC12 interface. The **long-buildout** statement is not supported for DS3 channels on a channelized OC12 interface.

On T3 interfaces, the default encapsulation is PPP.

For ATM1 interfaces, you can configure a subset of E3 options statements.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *T3 Interfaces Overview*

Copyright © 2018, Juniper Networks, Inc.

tag-protocol-id (TPIDs Expected to Be Sent or Received)

Syntax tag-protocol-id [*tpids*];

Hierarchy Level [edit interfaces *interface-name* gigether-options ethernet-switch-profile],
[edit interfaces *interface-name* aggregated-ether-options ethernet-switch-profile],
[edit interfaces *interface-name* aggregated-ether-options ethernet-switch-profile],
[edit interfaces *interface-name* ether-options ethernet-switch-profile]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.
Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
Statement introduced in Junos OS Release 14.1X53-D15 for the QFX Series.

Description For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, aggregated Ethernet with Gigabit Ethernet IQ interfaces, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC, and the built-in Gigabit Ethernet port on the M7i router), define the TPIDs expected to be sent or received on a particular VLAN. For each Gigabit Ethernet port, you can configure up to eight TPIDs using the **tag-protocol-id** statement; but only the first four TPIDs are supported on IQ2 and IQ2-E interfaces.

For 10-Gigabit Ethernet LAN/WAN PIC interfaces on T Series routers only the default TPID value (**0x8100**) is supported.

For Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces on EX Series switches, define the TPIDs expected to be sent or received on a particular VLAN. The default TPID value is **0x8100**. Other supported values are **0x88a8**, **0x9100**, and **0x9200**.

Options *tpids*—TPIDs to be accepted on the VLAN. Specify TPIDs in hexadecimal.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Frames with Particular TPIDs to Be Processed as Tagged Frames*
- *Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support (CLI Procedure)*

Copyright © 2018, Juniper Networks, Inc.

tag-protocol-id (TPID to Rewrite)

Syntax tag-protocol-id *tpid*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces only, configure the outer TPID value. All TPIDs you include in input and output VLAN maps must be among those you specify at the **[edit interfaces *interface-name* gigether-options ethernet-switch-profile tag-protocol-id [*tpids*]]** hierarchy level.

For 10-Gigabit Ethernet LAN/WAN PIC interfaces on T Series routers the default TPID value (**0x8100**) is supported.

Default If the **tag-protocol-id** statement is not configured, the TPID value is 0x8100.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Configuring Inner and Outer TPIDs and VLAN IDs*

targeted-broadcast

Syntax (EX Series, MX Series, ACX Series)

```
targeted-broadcast {
  forward-and-send-to-re;
  forward-only;
}
```

Syntax (QFX Series, OCX1100, EX4600, NFX Series)

```
targeted-broadcast;
```

Hierarchy Level (EX Series, MX Series, ACX Series)

[edit interfaces *interface-name* unit *logical-unit-number* family inet],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet]

Hierarchy Level (QFX Series, OCX1100, EX4600, NFX Series)

[edit interfaces *interface-name* unit *logical-unit-number* family inet],
[edit interfaces interface-range *interface-range-name* unit *logical-unit-number* family inet]

Release Information

Statement introduced in Junos OS Release 9.4 for EX Series switches.
Statement introduced in Junos OS Release 10.2.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

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 remaining statements are explained separately. See CLI Explorer.

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

- Configuring Targeted Broadcast on page 272
- Understanding Targeted Broadcast on page 271

 Copyright © 2018, Juniper Networks, Inc.

targeted-distribution (Static Interfaces over Aggregated Ethernet)

Syntax targeted-distribution;

Hierarchy Level [edit interfaces demux0 unit *logical-unit-number*],
[edit interfaces pp0 unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 11.2.
Statement introduced in Junos OS Release 13.2R2 for EX Series switches.

Description Configure egress data for a logical interface to be sent across a single member link in an aggregated Ethernet bundle. A backup link is provisioned and CoS scheduling resources are switched to the backup link in the event that the primary assigned link goes down. The aggregated Ethernet interface must be configured without link protection.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *CoS for PPPoE Subscriber Interfaces Overview*
- *Configuring the Distribution Type for PPPoE Subscribers on Aggregated Ethernet Interfaces*
- *Verifying the Distribution of PPPoE Subscribers in an Aggregated Ethernet Interface*
- *Targeted Traffic Distribution on Aggregated Ethernet Interfaces in a Virtual Chassis*
- *Configuring Module Redundancy for a Virtual Chassis*
- *Configuring Chassis Redundancy for a Virtual Chassis*

Copyright © 2018, Juniper Networks, Inc.

tei-option

Syntax tei-option (first-call | power-up);

Hierarchy Level [edit interfaces br-*pim*/0/*port*isdn-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For ISDN interfaces, configure when the Terminal Endpoint Identifier (TEI) negotiates with the ISDN provider.

Options **first-call**—Activation does not occur until the call setup is sent.

power-up—Activation occurs when the Services Router is powered on.

Default: power-up

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Junos OS Interfaces and Routing Configuration Guide*

 Copyright © 2018, Juniper Networks, Inc.

terminate (PPPoE Service Name Tables)

Syntax terminate;

Hierarchy Level [edit protocols pppoe service-name-tables *table-name* service *service-name*],
[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]

Release Information Statement introduced in Junos OS Release 10.0.
Support at **[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]** hierarchy level introduced in Junos OS Release 10.2.

Description Direct the router to immediately respond to a PPPoE Active Discovery Initiation (PADI) control packet received from a PPPoE client by sending the client a PPPoE Active Discovery Offer (PADO) packet. The PADO packet contains the name of the access concentrator (router) that can service the client request. The **terminate** action is the default action for a named service entry, **empty** service entry, **any** service entry, or agent circuit identifier/agent remote identifier (ACI/ARI) pair in a PPPoE service name table.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Configuring PPPoE Service Name Tables*

Copyright © 2018, Juniper Networks, Inc.

then

Syntax

```
then {
  discard;
}
```

Hierarchy Level

[edit firewall hierarchical-policer aggregate],
[edit firewall hierarchical-policer premium]

Release Information

Statement introduced in Junos OS Release 9.5.

Description

On M40e, M120, and M320 (with FFPC and SFPC) edge routers and T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs, discard packets when a specified bandwidth or burst limits for an aggregate level of a hierarchical policer is reached.

Options

discard—Discard packets if condition is met.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- Applying Policers on page 226
- *Class of Service Feature Guide for Routing Devices and EX9200 Switches*

Copyright © 2018, Juniper Networks, Inc.

threshold

Syntax threshold *bytes*;

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the bucket threshold, which controls the burstiness of the leaky bucket mechanism. The larger the value, the more bursty the traffic, which means that over a very short amount of time, the interface can receive or transmit close to line rate, but the average over a longer time is at the configured bucket rate.

Options ***bytes***—Maximum size, in bytes, for traffic bursts. For ease of entry, you can enter ***number*** either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000). For example, the entry **threshold 2k** corresponds to a threshold of 2000 bytes.

Range: 0 through 65,535 bytes

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Receive and Transmit Leaky Bucket Properties to Reduce Network Congestion on page 156

timeslots

Syntax timeslots *time-slot-range*;

Hierarchy Level [edit interfaces e1-*fpc/pic/port*],
[edit interfaces t1-*fpc/pic/port*],
[edit interfaces *interface-name* e1-options],
[edit interfaces *interface-name* partition *partition-number*],
[edit interfaces *interface-name* t1-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For E1 and T1interfaces, allocate the specific time slots by number.

NOTE: When configuring E1 or T1 interfaces on the 10-port Channelized E1/T1 IQE PIC, the **timeslots** statement must be included at the **[edit interfaces e1-*fpc/pic/port*]** or **[edit interfaces t1-*fpc/pic/port*]** hierarchy level as appropriate.

Options *time-slot-range*—Actual time slot numbers allocated:

Range: Ranges vary by interface type and configuration option as follows:

- 1 through 24 for T1 interfaces (0 is reserved)
- 1 through 31 for 4-port E1 PICs (0 is reserved)
- 1 through 31 for *Nx*DS0 interfaces (0 is reserved)
- 2 through 32 for 10-port Channelized E1 and 10-port Channelized E1 IQ PICs (1 is reserved)
- 2 through 32 for the setting under **e1-options** with IQE PICs (1 is reserved) (when creating fractional E1)
- 1 through 31 for the setting under **partition** with IQE PICs (0 is reserved) (when creating *Nx*DS0)

NOTE: When creating fractional E1 interfaces only, if you connect a 4-port E1 PIC interface to a device that uses time slot numbering from 2 through 32, you must subtract 1 from the configured number of time slots.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- *Configuring Fractional E1 IQ and IQE Interfaces*
- *Configuring Fractional T1 IQ and IQE Interfaces*
- *Configuring Fractional E1 Time Slots*
- *Configuring Fractional T1 Time Slots*
- *Configuring a Channelized T1/E1 Interface to Drop and Insert Time Slots*

tm

Syntax tm (ignore | normal | require);

Hierarchy Level [edit interfaces *interface-name* serial-options dce-options],
[edit interfaces *interface-name* serial-options dte-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For EIA-530 interfaces only, configure the from-DCE signal, test-mode (TM).

Options **ignore**—The from-DCE signal is ignored.

normal—Normal TM signal handling as defined by the TIA/EIA Standard 530.

require—The from-DCE signal must be asserted.

Default: normal

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Serial Signal Handling on page 344

tm-polarity

Syntax tm-polarity (negative | positive);

Hierarchy Level [edit interfaces *interface-name* serial-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure TM signal polarity.

Options **negative**—Negative signal polarity.

positive—Positive signal polarity.

Default: positive

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Serial Signal Polarities on page 347

 Copyright © 2018, Juniper Networks, Inc.

traceoptions (Individual Interfaces)

List of Syntax

Syntax (Individual interfaces with PTX Series, EX Series, ACX Series) on page 1019
Syntax (Individual interfaces with QFX Series, OCX1100, EX4600, NFX Series) on page 1019
Syntax (OAMLFM with EX Series, QFX Series, NFX Series) on page 1019
Syntax (Interface process with ACX Series, SRX Series, MX Series, M Series, T Series) on page 1019

Syntax (Individual interfaces with PTX Series, EX Series, ACX Series)

```
traceoptions {
  file filename <files name> <size size> <world-readable | no-world-readable>;
  flag flag;
  match;
}
```

Syntax (Individual interfaces with QFX Series, OCX1100, EX4600, NFX Series)

```
traceoptions {
  flag flag;
}
```

Syntax (OAMLFM with EX Series, QFX Series, NFX Series)

```
traceoptions {
  file filename <files number> <match regex> <size size> <world-readable |
    no-world-readable>;
  flag flag ;
  no-remote-trace;
}
```

Syntax (Interface process with ACX Series, SRX Series, MX Series, M Series, T Series)

```
traceoptions {
  file <filename> <files number> <match regular-expression> <size size> <world-readable |
    no-world-readable>;
  flag flag <disable>;
  no-remote-trace;
}
```

Hierarchy Level (Individual interfaces with PTX Series, EX Series, ACX Series, QFX Series, OCX1100, EX4600, NFX Series)

[edit interfaces *interface-name*]

Hierarchy Level (Interface process with ACX Series, SRX Series, MX Series, M Series, T Series)

[edit interfaces]

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 Metro Routers.

Copyright © 2018, Juniper Networks, Inc.

Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in JUNOS Release 10.2 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Define tracing operations for individual interfaces.

To specify more than one tracing operation, include multiple **flag** statements.

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/dcd**.

On EX Series, QFX Series, and NFX Series platforms, configure tracing options the link fault management.

On ACX Series, SRX Series, MX Series, M Series, and T Series platforms define tracing operations for the interface process (dcd).

Default If you do not include this statement, no interface-specific tracing operations are performed.

 Copyright © 2018, Juniper Networks, Inc.

Table 54: OptionOptionstraceoptions

Option	Individual interfaces with PTX Series, ACX Series, EX Series	Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series	Interface Process with OAMLFM with EX Series, QFX Series, NFX Series	Interface process with ACX Series, SRX Series, MX Series, M Series, T Series
file*filename*	—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log/dcd**. By default, interface process tracing output is placed in the file.	-	—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.	—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. By default, interface process tracing output is placed in the file **dcd**
files *number*	**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.	-	—(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 *x*k to specify KB, *x*m to specify MB, or *x*g to specify GB number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the size option.	—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option. *Range:* 2 through 1000 *Default:* 3 files
flag	—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. The following are the interface-specific tracing options. • **all**—All interface tracing operations • **event**—Interface events • **ipc**—Interface interprocess	—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. The following are the interface-specific tracing options. • **all**—All interface tracing operations • **event**—Interface events • **ipc**—Interface interprocess	—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags: • **action-profile**—Trace action profile invocation events. • **all**—Trace all events. • **configuration**—Trace configuration events. • **protocol**—Trace	—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags: • **all** • **change-events**—Log changes that produce configuration events • **config-states**—Log the configuration

Table 54: Options for traceoptions (continued)

Option	Individual interfaces with PTX Series, ACX Series, EX Series	Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series	Interface Process with OAMLFM with EX Series, QFX Series, NFX Series	Interface process with ACX Series, SRX Series, MX Series, M Series, T Series
	communication (IPC) messages • **media**—Interface media changes • **q921**—Trace ISDN Q.921 frames • **q931**—Trace ISDN Q.931 frames	communication (IPC) messages • **media**—Interface media changes • **q921**—Trace ISDN Q.921 frames • **q931**—Trace ISDN Q.931 frames	protocol processing events. • **routing socket**—Trace routing socket events.	state machine changes • **kernel**—Log configuration IPC messages to kernel • **kernel-detail**—Log details of configuration messages to kernel
match	—(Optional) Regular expression for lines to be traced.	-	—(Optional) Refine the output to log only those lines that match the given regular expression.	-
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.	-	—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the files option. *Syntax: x*k to specify KB, *x*m to specify MB, or *x*g to specify GB *Range:* 10 KB through 1 GB *Default:* 128 KB *Default:* If you do not include this option, tracing output is appended to an existing trace file.	

 Copyright © 2018, Juniper Networks, Inc.

Table 54: Options for traceoptions (continued)

Option	Individual interfaces with PTX Series, ACX Series, EX Series	Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series	Interface Process with OAMLFM with EX Series, QFX Series, NFX Series	Interface process with ACX Series, SRX Series, MX Series, M Series, T Series
				—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option. *Syntax:* ***x*k** to specify kilobytes, ***x*m** to specify megabytes, or ***x*g** to specify gigabytes *Range:* 10 KB through the maximum file size supported on your router *Default:* 1 MB
no-world-readable	—(Optional) Prevent any user from reading the log file.	-	—(Optional) Restrict file access to the user who created the file.	—(Optional) Disallow any user to read the log file.
world-readable	—(Optional) Allow any user to read the log file.	-	—(Optional) Enable unrestricted file access.	—(Optional) Allow any user to read the log file.
disable	-	-	-	

Table 54: Options for traceoptions (continued)

Option	Individual interfaces with PTX Series, ACX Series, EX Series	Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series	Interface Process with OAMLFM with EX Series, QFX Series, NFX Series	Interface process with ACX Series, SRX Series, MX Series, M Series, T Series
				—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.
no-remote-trace	-	-	—(Optional) Disable the remote trace.	-
match ***regex***	-	-	-	—(Optional) Refine the output to include only those lines that match the given regular expression.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Tracing Operations of an Individual Router Interface on page 355
- *Tracing Operations of an Individual Router or Switch Interface*
- *Example: Configuring Ethernet OAM Link Fault Management*
- *Configuring Ethernet OAM Link Fault Management (CLI Procedure)*
- Tracing Operations of the Interface Process on page 356

 Copyright © 2018, Juniper Networks, Inc.

traceoptions (Individual Interfaces)

List of Syntax

Syntax (Individual interfaces with PTX Series, EX Series, ACX Series) on page 1025
Syntax (Individual interfaces with QFX Series, OCX1100, EX4600, NFX Series) on page 1025
Syntax (OAMLFM with EX Series, QFX Series, NFX Series) on page 1025
Syntax (Interface process with ACX Series, SRX Series, MX Series, M Series, T Series) on page 1025

Syntax (Individual interfaces with PTX Series, EX Series, ACX Series)

```
traceoptions {
  file filename <files name> <size size> <world-readable | no-world-readable>;
  flag flag;
  match;
}
```

Syntax (Individual interfaces with QFX Series, OCX1100, EX4600, NFX Series)

```
traceoptions {
  flag flag;
}
```

Syntax (OAMLFM with EX Series, QFX Series, NFX Series)

```
traceoptions {
  file filename <files number> <match regex> <size size> <world-readable |
    no-world-readable>;
  flag flag ;
  no-remote-trace;
}
```

Syntax (Interface process with ACX Series, SRX Series, MX Series, M Series, T Series)

```
traceoptions {
  file <filename> <files number> <match regular-expression> <size size> <world-readable |
    no-world-readable>;
  flag flag <disable>;
  no-remote-trace;
}
```

Hierarchy Level (Individual interfaces with PTX Series, EX Series, ACX Series, QFX Series, OCX1100, EX4600, NFX Series)

[edit interfaces *interface-name*]

Hierarchy Level (Interface process with ACX Series, SRX Series, MX Series, M Series, T Series)

[edit interfaces]

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 Metro Routers.

Copyright © 2018, Juniper Networks, Inc.

Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in JUNOS Release 10.2 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Define tracing operations for individual interfaces.

To specify more than one tracing operation, include multiple **flag** statements.

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/dcd**.

On EX Series, QFX Series, and NFX Series platforms, configure tracing options the link fault management.

On ACX Series, SRX Series, MX Series, M Series, and T Series platforms define tracing operations for the interface process (dcd).

Default If you do not include this statement, no interface-specific tracing operations are performed.

Copyright © 2018, Juniper Networks, Inc.

Table 55: OptioOptionsfor traceoptions

Option	Individual interfaces with PTX Series, ACX Series, EX Series	Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series	Interface Process with OAMLFM with EX Series, QFX Series, NFX Series	Interface process with ACX Series, SRX Series, MX Series, M Series, T Series
file*filename*	—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log/dcd**. By default, interface process tracing output is placed in the file.	-	—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.	—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. By default, interface process tracing output is placed in the file **dcd**
files *number*	**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.	-	—(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 *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the size option.	—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option. *Range:* 2 through 1000 *Default:* 3 files
flag	—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. The following are the interface-specific tracing options. • **all**—All interface tracing operations • **event**—Interface events • **ipc**—Interface interprocess	—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. The following are the interface-specific tracing options. • **all**—All interface tracing operations • **event**—Interface events • **ipc**—Interface interprocess	—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags: • **action-profile**—Trace action profile invocation events. • **all**—Trace all events. • **configuration**—Trace configuration events. • **protocol**—Trace	—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags: • **all** • **change-events**—Log changes that produce configuration events • **config-states**—Log the configuration

Table 55: Options for traceoptions (continued)

Option	Individual interfaces with PTX Series, ACX Series, EX Series	Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series	Interface Process with OAMLFM with EX Series, QFX Series, NFX Series	Interface process with ACX Series, SRX Series, MX Series, M Series, T Series
	communication (IPC) messages • **media**—Interface media changes • **q921**—Trace ISDN Q.921 frames • **q931**—Trace ISDN Q.931 frames	communication (IPC) messages • **media**—Interface media changes • **q921**—Trace ISDN Q.921 frames • **q931**—Trace ISDN Q.931 frames	protocol processing events. • **routing socket**—Trace routing socket events.	state machine changes • **kernel**—Log configuration IPC messages to kernel • **kernel-detail**—Log details of configuration messages to kernel
match	—(Optional) Regular expression for lines to be traced.	-	—(Optional) Refine the output to log only those lines that match the given regular expression.	-
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.	-	—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the files option. *Syntax:* *x*k to specify KB, *x*m to specify MB, or *x*g to specify GB *Range:* 10 KB through 1 GB *Default:* 128 KB *Default:* If you do not include this option, tracing output is appended to an existing trace file.	

 Copyright © 2018, Juniper Networks, Inc.

Table 55: Options for traceoptions (continued)

Option	Individual interfaces with PTX Series, ACX Series, EX Series	Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series	Interface Process with OAMLFM with EX Series, QFX Series, NFX Series	Interface process with ACX Series, SRX Series, MX Series, M Series, T Series
				—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option. *Syntax:* ***x*k** to specify kilobytes, ***x*m** to specify megabytes, or ***x*g** to specify gigabytes *Range:* 10 KB through the maximum file size supported on your router *Default:* 1 MB
no-world-readable	—(Optional) Prevent any user from reading the log file.	-	—(Optional) Restrict file access to the user who created the file.	—(Optional) Disallow any user to read the log file.
world-readable	—(Optional) Allow any user to read the log file.	-	—(Optional) Enable unrestricted file access.	—(Optional) Allow any user to read the log file.
disable	-	-	-	

Table 55: Options for traceoptions (continued)

Option	Individual interfaces with PTX Series, ACX Series, EX Series	Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series	Interface Process with OAMLFM with EX Series, QFX Series, NFX Series	Interface process with ACX Series, SRX Series, MX Series, M Series, T Series
				—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.
no-remote-trace	-	-	—(Optional) Disable the remote trace.	-
match *regex*	-	-	-	—(Optional) Refine the output to include only those lines that match the given regular expression.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Tracing Operations of an Individual Router Interface on page 355
- *Tracing Operations of an Individual Router or Switch Interface*
- *Example: Configuring Ethernet OAM Link Fault Management*
- *Configuring Ethernet OAM Link Fault Management (CLI Procedure)*
- Tracing Operations of the Interface Process on page 356

Copyright © 2018, Juniper Networks, Inc.

traceoptions (LACP)

Syntax

traceoptions {
  file <*filename*> <files *number*> <size *size*> <world-readable | no-world-readable>;
  flag *flag*;
  no-remote-trace;
}

Hierarchy Level [edit protocols lacp]

Release Information Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.

Description Define tracing operations for the LACP protocol.

Default If you do not include this statement, no LACP protocol tracing operations are performed.

Options *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, interface process tracing output is placed in the file **lacpd**.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—All LACP tracing operations
- **configuration**—Configuration code
- **packet**—Packets sent and received
- **process**—LACP process events
- **protocol**—LACP protocol state machine
- **routing-socket**—Routing socket events
- **startup**—Process startup events

no-world-readable—(Optional) Prevent any user from reading the log file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0.** When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option:

Syntax: ***x*k** to specify kilobytes, ***x*m** to specify megabytes, or ***x*g** to specify gigabytes

Range: 10 KB through the maximum file size supported on your router

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Tracing LACP Operations*

 Copyright © 2018, Juniper Networks, Inc.

traceoptions (PPP Process)

Syntax

```
traceoptions {
  file filename <files number> <match regular-expression> <size size> <world-readable |
    no-world-readable>;
  flag flag;
  level severity-level;
  no-remote-trace;
}
```

Hierarchy Level [edit protocols ppp]

Release Information Statement introduced in Junos OS Release 7.5.

Description Define tracing operations for the PPP process.

To specify more than one tracing operation, include multiple **flag** statements.

You cannot specify a separate trace tile. Tracing information is placed in the system **syslog** file in the directory **/var/log/pppd**.

Default If you do not include this statement, no PPPD-specific tracing operations are performed.

Options ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, commit script process tracing output is placed in the file **ppd**. If you include the **file** statement, you must specify a filename. To retain the default, you can specify **eventd** as the filename.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

Range: 2 through 1000

Default: 3 files

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. The following are the PPPD-specific tracing options.

- **access**—Access code
- **address-pool**—Address pool code

- **all**—All areas of code
- **auth**—Authentication code
- **chap**—Challenge Handshake Authentication Protocol (CHAP) code
- **config**—Configuration code
- **ifdb**—Interface database code
- **lcp**—LCP state machine code
- **memory**—Memory management code
- **message**—Message processing code
- **mlppp**—Trace MLPPP code
- **ncp**—NCP state machine code
- **pap**—Password Authentication Protocol (PAP) code
- **ppp**—PPP protocol processing code
- **radius**—RADIUS processing code
- **rtsock**—Routing socket code
- **session**—Session management code
- **signal**—Signal handling code
- **timer**—Timer code
- **ui**—User interface code

match *regex*—(Optional) Refine the output to include only those lines that match the given regular expression.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and filename.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through 1 GB

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

non-world-readable—(Optional) By default, log files can be accessed only by the user who configures the tracing operation. Specify **non-world-readable** to reset the default.

 Copyright © 2018, Juniper Networks, Inc.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Tracing Operations of the pppd Process on page 153

traceoptions (PPPoE)

Syntax

```
traceoptions {
  file <filename> <files number> <match regular-expression > <size maximum-file-size>
    <world-readable | no-world-readable>;
  filter {
    aci regular-expression;
    ari regular-expresion;
    service-name regular-expresion;
    underlying-interface interface-name;
  }
  flag flag;
  level (all | error | info | notice | verbose | warning);
  no-remote-trace;
}
```

Hierarchy Level

[edit protocols pppoe]

Release Information

Statement introduced in Junos OS Release 9.6.
Option **filter** introduced in Junos OS Release 12.3

Description

Define tracing operations for PPPoE processes.

Options

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files *number*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

disable—Disable this trace flag.

filter—Additional filter to refine the output to display particular subscribers. Filtering based on the following subscriber identifiers simplifies troubleshooting in a scaled environment.

BEST PRACTICE: **Due to the complexity of agent circuit identifiers and agent remote identifiers, we recommend that you do not try an exact match when filtering on these options. For service names, searching on the exact name is appropriate, but you can also use a regular expression with that option.**

- **aci *regular-expression***—Regular expression to match the agent circuit identifier provided by PPPoE client.

 Copyright © 2018, Juniper Networks, Inc.

- **ari** *regular-expression*—Regular expression to match the agent remote identifier provided by PPPoE client.
- **service** *regular-expression*—Regular expression to match the name of PPPoE service.
- **underlying-interface** *interface-name*—Name of a PPPoE underlying interface. You cannot use a regular expression for this filter option.

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—Trace all operations.
- **config**—Trace configuration events.
- **events**—Trace events.
- **gres**—Trace GRES events.
- **init**—Trace initialization events.
- **interface-db**—Trace interface database operations.
- **memory**—Trace memory processing events.
- **protocol**—Trace protocol events.
- **rtsock**—Trace routing socket events.
- **session-db**—Trace connection events and flow.
- **signal**—Trace signal operations.
- **state**—Trace state handling events.
- **timer**—Trace timer processing.
- **ui**—Trace user interface processing.

level—Level of tracing to perform. You can specify any of the following levels:

- **all**—Match all levels.
- **error**—Match error conditions.
- **info**—Match informational messages.
- **notice**—Match notice messages about conditions requiring special handling.
- **verbose**—Match verbose messages.
- **warning**—Match warning messages.

Default: **error**

match *regular-expression*—(Optional) Refine the output to include lines that contain the regular expression.

no-remote-trace—Disable remote tracing.

no-world-readable—(Optional) Disable unrestricted file access.

size ***maximum-file-size***—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: ***size*****k** to specify KB, ***size*****m** to specify MB, or ***size*****g** to specify GB

Range: 10240 through 1073741824

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level
trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation
- *Configuring PPPoE Service Name Tables*
- *Tracing PPPoE Operations*

translate-discard-eligible

Syntax (translate-discard-eligible | no-translate-discard-eligible);

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family ccc],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family ccc]

Release Information Statement introduced before Junos OS Release 7.4.

Description For interfaces with encapsulation type Frame Relay CCC, enable or disable translation of Frame Relay discard eligible (DE) control bits.

Default DE bit translation is disabled.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Configuring Frame Relay Control Bit Translation*

 Copyright © 2018, Juniper Networks, Inc.

translate-fecn-and-becn

Syntax (translate-fecn-and-becn | no-translate-fecn-and-becn);

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family ccc],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family ccc]

Release Information Statement introduced before Junos OS Release 7.4.

Description For interfaces with encapsulation type Frame Relay CCC, enable or disable translation of Frame Relay forward explicit congestion notification (FECN) control bits and Frame Relay backward explicit congestion notification (BECN) control bits.

Default FECN and BECN bit translation is disabled.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Frame Relay Control Bit Translation*

translate-plp-control-word-de

Syntax translate-plp-control-word-de

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family ccc]

Release Information Statement introduced before Junos OS Release 11.1.

Description For the interfaces with encapsulation type Frame Relay CCC, classify and rewrite the control word discard eligibility (DE) bit based on the packet loss priority (PLP).

Default PLP bit translation is disabled.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Frame Relay Control Bit Translation*
- *frame-relay-de*

transmit-bucket

Syntax

```
transmit-bucket {
  overflow discard;
  rate percentage;
  threshold bytes;
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Set parameters for the transmit leaky bucket, which specifies what percentage of the interface's total capacity can be used to transmit packets.

For each DS3 channel in a channelized OC12 interface, you can configure a unique transmit bucket.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Receive and Transmit Leaky Bucket Properties to Reduce Network Congestion on page 156
- **receive-bucket on page 921**

transmit-clock

Syntax transmit-clock invert;

Hierarchy Level [edit interfaces *interface-name* serial-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the transmit clock signal.

Options **invert**—Shift the clock phase 180 degrees.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Serial Clocking Mode on page 342

 Copyright © 2018, Juniper Networks, Inc.

transmit-period

Syntax transmit-period *seconds*;

Hierarchy Level [edit protocols dot1x authenticator interface *interface-id*]

Release Information Statement introduced in Junos OS Release 9.3.

Description Set the number of seconds the port waits before retransmitting the initial EAPOL PDUs to the client.

Options ***seconds***—The number of seconds the port waits before retransmitting the initial EAPOL PDUs to the client.

Range: 1 through 65,535 seconds

Default: 30 seconds

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *IEEE 802.1x Port-Based Network Access Control Overview*
- authenticator on page 448
- dot1x on page 543
- interface (IEEE 802.1x) on page 682

Copyright © 2018, Juniper Networks, Inc.

transmit-weight (ATM2 IQ CoS Forwarding Class)

Syntax transmit-weight (cells *number* | percent *number*);

Hierarchy Level [edit interfaces *interface-name* atm-options scheduler-maps *map-name* forwarding-class *class-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For ATM2 IQ interfaces only, assign a transmission weight to a forwarding class.

Default 95 percent for queue 0, 5 percent for queue 3.

Options **percent *percent***—Transmission weight of the forwarding class as a percentage of the total bandwidth.

Range: 5 through 100

cells *number*—Transmission weight of the forwarding class as a number of cells.

Range: 0 through 32,000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *ATM2 IQ VC Tunnel CoS Components Overview*

 Copyright © 2018, Juniper Networks, Inc.

transmit-weight (ATM2 IQ Virtual Circuit)

Syntax
transmit-weight *number*;

Hierarchy Level
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information
Statement introduced before Junos OS Release 7.4.

Description
For ATM2 IQ PICs only, configure the transmission weight.

Each VC is serviced in weighted round robin (WRR) mode. When VCs have data to send, they send the number of cells equal to their weight before passing control to the next active VC. This allows proportional bandwidth sharing between multiple VCs within a rate-shaped VP tunnel. VP tunnels are not supported on multipoint interfaces.

Options
number—Number of cells a VC sends before passing control to the next active VC within a VP tunnel.

Range: 1 through 32,767

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Configuring the ATM2 IQ Transmission Weight*

traps

Syntax (traps | no-traps);

Hierarchy Level (ACX Series, MX Series, T Series, M Series, SRX Series, EX Series)

[edit dynamic-profiles *profile-name* interfaces *interface-name*],
[edit interfaces *interface-name*],
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit interfaces interface-range *name*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Hierarchy Level (QFX Series, EX4600)

[edit interfaces *interface-name*],
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit interfaces interface-range *interface-range-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.
Support at the **[edit dynamic-profiles *profile-name* interfaces*interface-name*]** hierarchy level introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

Description

Enable or disable the sending of Simple Network Management Protocol (SNMP) notifications when the state of the connection changes.

(Enhanced subscriber management for MX Series routers) To enable SNMP notifications, you must first configure the **interface-mib** statement at the **[edit dynamic-profiles *profile-name* interfaces *interface-name*]** hierarchy level. If **interface-mib** is not configured, the **traps** statement has no effect.

BEST PRACTICE: **To achieve maximum performance when enhanced subscriber management is enabled, we recommend that you *not* enable SNMP notifications on all dynamic subscriber interfaces.**

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Enabling or Disabling SNMP Notifications on Physical Interfaces on page 171
- Enabling or Disabling SNMP Notifications on Logical Interfaces on page 200

 Copyright © 2018, Juniper Networks, Inc.

trigger

Syntax

```
trigger {
  defect ignore;
  defect hold-time up milliseconds down milliseconds;
}
```

Hierarchy Level [edit interfaces *interface-name* sonet-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For ATM over SONET/SDH, SONET/SDH interfaces, and 10-Gigabit Ethernet interfaces in WAN PHY mode, configure SONET/SDH defect triggers to be ignored.

Default If you do not include this statement, all SONET/SDH defect triggers are honored.

Options ***defect***—Defect to ignore or hold. It can be one of the following:

- **ais-l**—Line alarm indication signal
- **ais-p**—Path alarm indication signal
- **ber-sd**—Bit error rate signal degrade
- **ber-sf**—Bit error rate signal fault
- **locd** (ATM only)—Loss of cell delineation
- **lof**—Loss of frame
- **lol**—PHY loss of light
- **lop-p**—Path loss of pointer
- **los**—Loss of signal
- **pll**—PHY phase-locked loop out of lock
- **plm-p**—Path payload (signal) label mismatch
- **rfi-l**—Line remote failure indication
- **rfi-p**—Path remote failure indication
- **uneq-p**—Path unequipped

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring SONET/SDH Defect Triggers*

trigger-link-failure

Syntax [trigger-link-failure *interface-name*];

Hierarchy Level [edit interfaces lsq-*fpc/pic/port* lsq-failure-options]

Release Information Statement introduced in Junos OS Release 7.4.

Description List of SONET interfaces connected to the LSQ interface that can implement Automatic Protection Switching (APS) if the LSQ PIC fails.

Options ***interface-name***—Name of SONET interface.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Junos OS Services Interfaces Library for Routing Devices*

 Copyright © 2018, Juniper Networks, Inc.

trunk-bandwidth

Syntax trunk-bandwidth *rate*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For ATM2 IQ interfaces configured to use Layer 2 circuit trunk mode, configure a scheduler so that unused bandwidth from any inactive trunk is proportionally shared among the active trunks.

During congestion, each trunk receives a proportional share of the leftover bandwidth, thus minimizing the latency on each trunk.

Options *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.

Range: 1,000,000 through 542,526,792 bps

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Layer 2 Circuit Trunk Mode Scheduling Overview*

Copyright © 2018, Juniper Networks, Inc.

trunk-id

Syntax trunk-id *number*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For ATM2 IQ interfaces with ATM CCC cell-relay encapsulation, configure the trunk identification number.

When you associate a trunk ID number with a logical interface, you are in effect specifying the interfaces that are allowed to send ATM traffic over an LSP.

Options ***number***—A valid trunk identifier.

Range: For UNI mode, 0 through 7. For NNI mode, 0 through 31.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Configuring Layer 2 Circuit Transport Mode*

ttl

Syntax ttl *value*;

Hierarchy Level [edit interfaces *interface-name* unit *number* tunnel]

Release Information Statement introduced before Junos OS Release 7.4

Description Set the time-to-live value bit in the header of the outer IP packet.

Options ***value***—Time-to-live value.

Range: 0 through 255

Default: 64

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Junos OS Services Interfaces Library for Routing Devices*

 Copyright © 2018, Juniper Networks, Inc.

tunnel

Syntax

```
tunnel {
  backup-destination address;
  destination address;
  key number;
  routing-instance {
    destination routing-instance-name;
  }
  source source-address;
  ttl number;
}
```

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure a tunnel. You can use the tunnel for unicast and multicast traffic or just for multicast traffic. You can also use tunnels for encrypted traffic or VPNs.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*
- *Junos OS VPNs Library for Routing Devices*

underlying-interface

Syntax underlying-interface *interface-name*;

Hierarchy Level [edit interfaces pp0 unit *logical-unit-number* pppoe-options],
[edit interfaces demux0 unit *logical-unit-number* demux-options],
[edit logical-systems *logical-system-name* interfaces demux0 unit *logical-unit-number* demux-options],
[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number* pppoe-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* interfaces demux0 unit *logical-unit-number* demux-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* interfaces pp0 unit *logical-unit-number* pppoe-options]

Release Information Statement introduced before Junos OS Release 7.4.
Support for aggregated Ethernet added in Junos OS Release 9.4.

Description Configure the interface on which PPP over Ethernet is running.

For demux interfaces, configure the underlying interface on which the demultiplexing (demux) interface is running.

Options *interface-name*—Name of the interface on which PPP over Ethernet or demux is running. For example, **at-0/0/1.0** (ATM VC), **fe-1/0/1.0** (Fast Ethernet interface), **ge-2/0/0.0** (Gigabit Ethernet interface), **ae1.0** (for IP demux on an aggregated Ethernet interface), or **ae1** (for VLAN demux on an aggregated Ethernet interface).

NOTE: Demux interfaces are currently supported on Gigabit Ethernet, Fast Ethernet, 10-Gigabit Ethernet interfaces, or aggregated Ethernet devices.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring an IP Demultiplexing Interface on page 316
- Configuring a VLAN Demultiplexing Interface on page 320
- *Configuring the PPPoE Underlying Interface*
- *Junos OS Interfaces and Routing Configuration Guide*

Copyright © 2018, Juniper Networks, Inc.

unframed

Syntax (unframed | no-unframed);

Hierarchy Level [edit interfaces *interface-name* e3-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For E3 IQ interfaces only, enable or disable unframed mode. In unframed mode, the E3 IQ interface do not detect yellow (**ylw**) or loss-of-frame (**lof**) alarms.

Default Unframed mode is disabled.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring E3 IQ and IQE Unframed Mode*

unidirectional

Syntax unidirectional;

Hierarchy Level [edit interfaces *interface-name*],
[edit logical-systems *logical-system-name* interfaces *interface-name*]

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:

- 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

- Understanding Unidirectional Traffic Flow on Physical Interfaces on page 157
- Enabling Unidirectional Traffic Flow on Physical Interfaces on page 158

unit (Dynamic Profiles Standard Interface)

Syntax

```
unit logical-unit-number {
  actual-transit-statistics;
  auto-configure {
    agent-circuit-identifier {
      dynamic-profile profile-name;
    }
    line-identity {
      include {
        accept-no-ids;
        circuit-id;
        remote-id;
      }
      dynamic-profile profile-name;
    }
  }
  dial-options {
    ipsec-interface-id name;
    l2tp-interface-id name;
    (shared | dedicated);
  }
  encapsulation (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-cisco-nlpid | atm-tcc-vc-mux
    | atm-mlppp-llc | atm-nlpid | atm-ppp-llc | atm-ppp-vc-mux | atm-snap | atm-tcc-snap
    | atm-vc-mux | ether-over-atm-llc | ether-vpls-over-atm-llc | ether-vpls-over-fr |
    ether-vpls-over-ppp | ethernet | frame-relay-ccc | frame-relay-ppp | frame-relay-tcc |
    frame-relay-ether-type | frame-relay-ether-type-tcc | multilink-frame-relay-end-to-end
    | multilink-ppp | ppp-over-ether | ppp-over-ether-over-atm-llc | vlan-bridge | vlan-ccc |
    vlan-vci-ccc | vlan-tcc | vlan-vpls);
  family family {
    address address;
    demux-destination,
    filter {
      adf {
        counter;
        input-precedence precedence;
        not-mandatory;
        output-precedence precedence;
        rule rule-value;
      }
      input filter-name {
        precedence precedence;
        shared-name filter-shared-name;
      }
      output filter-name {
        precedence precedence;
        shared-name filter-shared-name;
      }
    }
    max-sessions number;
    max-sessions-vsa-ignore;
    rpf-check {
      fail-filter filter-name;
      mode loose;
```


Copyright © 2018, Juniper Networks, Inc.

```
  }
  service {
    input {
      service-set service-set-name {
        service-filter filter-name;
      }
      post-service-filter filter-name;
    }
    input-vlan-map {
      inner-tag-protocol-id tpid;
      inner-vlan-id number;
      (push | swap);
      tag-protocol-id tpid;
      vlan-id number;
    }
    output {
      service-set service-set-name {
        service-filter filter-name;
      }
    }
    output-vlan-map {
      inner-tag-protocol-id tpid;
      inner-vlan-id number;
      (pop | swap);
      tag-protocol-id tpid;
      vlan-id number;
    }
  }
  service-name-table table-name
  short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
    maximum-seconds>;
  unnumbered-address interface-name <preferred-source-address address>;
}
keepalives {
  interval seconds;
  }
ppp-options {
  aaa-options aaa-options-name;
  authentication [ authentication-protocols ];
  chap {
    challenge-length minimum minimum-length maximum maximum-length;
    local-name name;
  }
  ignore-magic-number-mismatch;
  initiate-ncp (dual-stack-passive | ipv6 | ip)
  ipcp-suggest-dns-option;
  mru size;
  mtu (size | use-lower-layer);
  on-demand-ip-address;
  pap;
  peer-ip-address-optional;
  local-authentication {
    password password;
    username-include {
      circuit-id;
      delimiter character;
```

```
                domain-name name;
                mac-address;
                remote-id;
              }
            }
          }
          vlan-id number;
          vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
          filter {
            input filter-name {
              shared-name  filter-shared-name;
            }
            output filter-name {
              shared-name  filter-shared-name;
            }
          }
          host-prefix-only;
          service {
            pcef  pcef-profile-name {
              activate rule-name | activate-all;
            }
          }
        }
```

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options *logical-unit-number*—The specific unit number of the interface you want to assign to the dynamic profile, or one of the following predefined variables:

- **$junos-underlying-interface-unit**—For static VLANs, the unit number variable. The static unit number variable is dynamically replaced with the client unit number when the client session begins. The client unit number is specified by the DHCP when it accesses the subscriber network.
- **$junos-interface-unit**—The unit number variable on a dynamic underlying VLAN interface for which you want to enable the creation of dynamic VLAN subscriber interfaces based on the ACI.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- *Configuring Dynamic Underlying VLAN Interfaces to Use Agent Circuit Identifier Information*
- *Configuring Static Underlying VLAN Interfaces to Use Agent Circuit Identifier Information*
- *Agent Circuit Identifier-Based Dynamic VLANs Overview*

unit

Syntax

```
unit logical-unit-number {
  accept-source-mac {
    mac-address mac-address {
      policer {
        input cos-policer-name;
        output cos-policer-name;
      }
    }
  }
  accounting-profile name;
  advisory-options {
    downstream-rate rate;
    upstream-rate rate;
  }
  allow-any-vci;
  atm-scheduler-map (map-name | default);
  auto-configure {
    agent-circuit-identifier {
      dynamic-profile profile-name;
    }
    line-identity {
      include {
        accept-no-ids;
        circuit-id;
        remote-id;
      }
      dynamic-profile profile-name;
    }
  }
  backup-options {
    interface interface-name;
  }
  bandwidth rate;
  cell-bundle-size cells;
  clear-dont-fragment-bit;
  compression {
    rtp {
      maximum-contexts number <force>;
      f-max-period number;
      queues [queue-numbers];
      port {
        minimum port-number;
        maximum port-number;
      }
    }
  }
  compression-device interface-name;
  copy-tos-to-outer-ip-header;
  demux-destination family;
  demux-source family;
  demux-options {
    underlying-interface interface-name;
```

 Copyright © 2018, Juniper Networks, Inc.

```
}
description text;
etree-ac-role (leaf | root);
interface {
  l2tp-interface-id name;
  (dedicated | shared);
}
dialer-options {
  activation-delay seconds;
  callback;
  callback-wait-period time;
  deactivation-delay seconds;
  dial-string [dial-string-numbers];
  idle-timeout seconds;
  incoming-map {
    caller caller-id | accept-all;
    initial-route-check seconds;
    load-interval seconds;
    load-threshold percent;
    pool pool-name;
    redial-delay time;
    watch-list {
      [routes];
    }
  }
}
disable;
disable-mlppp-inner-ppp-pfc;
dlci dlci-identifier;
drop-timeout milliseconds;
dynamic-call-admission-control {
  activation-priority priority;
  bearer-bandwidth-limit kilobits-per-second;
}
encapsulation type;
epd-threshold cells plp1 cells;
family family-name {
  ... the family subhierarchy appears after the main [edit interfaces interface-name unit
    logical-unit-number] hierarchy ...
}
fragment-threshold bytes;
host-prefix-only;
inner-vlan-id-range start start-id end end-id;
input-vlan-map {
  (pop | pop-pop | pop-swap | push | push-push | swap |
  swap-push | swap-swap);
  inner-tag-protocol-id tpid;
  inner-vlan-id number;
  tag-protocol-id tpid;
  vlan-id number;
}
interleave-fragments;
inverse-arp;
layer2-policer {
  input-policer policer-name;
  input-three-color policer-name;
```

```
    output-policer policer-name;
    output-three-color policer-name;
}
link-layer-overhead percent;
minimum-links number;
mrru bytes;
multicast-dlci dlci-identifier;
multicast-vci vpi-identifier.vci-identifier;
multilink-max-classes number;
multipoint;
oam-liveness {
    up-count cells;
    down-count cells;
}
oam-period (disable | seconds);
output-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap |
    swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
}
passive-monitor-mode;
peer-unit unit-number;
plp-to-clp;
point-to-point;
ppp-options {
    mru size;
    mtu (size | use-lower-layer);
    chap {
        access-profile name;
        default-chap-secret name;
        local-name name;
        passive;
    }
    compression {
        acfc;
        pfc;
    }
    dynamic-profile profile-name;
    ipcp-suggest-dns-option;
    lcp-restart-timer milliseconds;
    loopback-clear-timer seconds;
    ncp-restart-timer milliseconds;
    pap {
        access-profile name;
        default-pap-password password;
        local-name name;
        local-password password;
        passive;
    }
}
pppoe-options {
    access-concentrator name;
    auto-reconnect seconds;
    (client | server);
```

 Copyright © 2018, Juniper Networks, Inc.

```
    service-name name;
    underlying-interface interface-name;
  }
  pppoe-underlying-options {
    access-concentrator name;
    direct-connect;
    dynamic-profile profile-name;
    max-sessions number;
  }
  proxy-arp;
  service-domain (inside | outside);
  shaping {
    (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate burst
      length);
    queue-length number;
  }
  short-sequence;
  targeted-distribution;
  transmit-weight number;
  (traps | no-traps);
  trunk-bandwidth rate;
  trunk-id number;
  tunnel {
    backup-destination address;
    destination address;
    key number;
    routing-instance {
      destination routing-instance-name;
    }
    source source-address;
    ttl number;
  }
  vci vpi-identifier.vci-identifier;
  vci-range start start-vci end end-vci;
  vpi vpi-identifier;
  vlan-id number;
  vlan-id-range number-number;
  vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
  family family {
    accounting {
      destination-class-usage;
      source-class-usage {
        (input | output | input output);
      }
    }
    access-concentrator name;
    address address {
      ... the address subhierarchy appears after the main [edit interfaces interface-name unit
        logical-unit-number family family-name] hierarchy ...
    }
    bundle interface-name;
    core-facing;
    demux-destination {
      destination-prefix;
    }
    demux-source {
```

```
    source-prefix;
  }
  direct-connect;
  duplicate-protection;
  dynamic-profile profile-name;
  filter {
    group filter-group-number;
    input filter-name;
    input-list [filter-names];
    output filter-name;
    output-list [filter-names];
  }
  interface-mode (access | trunk);
  ipsec-sa sa-name;
  keep-address-and-control;
  mac-validate (loose | strict);
  max-sessions number;
  mtu bytes;
  multicast-only;
  no-redirects;
  policer {
    arp policer-template-name;
    input policer-template-name;
    output policer-template-name;
  }
  primary;
  protocols [inet iso mpls];
  proxy inet-address address;
  receive-options-packets;
  receive-ttl-exceeded;
  remote (inet-address address | mac-address address);
  rpf-check {
    fail-filter filter-name
    mode loose;
  }
  sampling {
    input;
    output;
  }
  service {
    input {
      post-service-filter filter-name;
      service-set service-set-name <service-filter filter-name>;
    }
    output {
      service-set service-set-name <service-filter filter-name>;
    }
  }
  service-name-table table-name
  (translate-discard-eligible | no-translate-discard-eligible);
  (translate-fecn-and-becn | no-translate-fecn-and-becn);
  translate-plp-control-word-de;
  unnumbered-address interface-name destination address
    destination-profile profile-name;
  vlan-id number;
  vlan-id-list [number number-number];
```

 Copyright © 2018, Juniper Networks, Inc.

```
                address address {
                  arp ip-address (mac | multicast-mac) mac-address <publish>;
                  broadcast address;
                  destination address;
                  destination-profile name;
                  eui-64;
                  master-only;
                  multipoint-destination address {
                    dlci dlci-identifier;
                    epd-threshold cells <plp1 cells>;
                    inverse-arp;
                    oam-liveness {
                      up-count cells;
                      down-count cells;
                    }
                    oam-period (disable | seconds);
                    shaping {
                      (cbr rate | rtvbr burst length peak rate sustained rate | vbr burst length peak rate
                        sustained rate);
                      queue-length number;
                    }
                    vci vpi-identifier.vci-identifier;
                  }
                  preferred;
                  primary;
                  (vrrp-group | vrrp-inet6-group) group-number {
                    (accept-data | no-accept-data);
                    advertise–interval seconds;
                    authentication-type authentication;
                    authentication-key key;
                    fast-interval milliseconds;
                    (preempt | no-preempt) {
                      hold-time seconds;
                    }
                    priority number;
                    track {
                      interface interface-name {
                        bandwidth-threshold bits-per-second priority-cost number;
                      }
                      priority-hold-time seconds;
                      route ip-address/prefix-length routing-instance instance-name priority-cost cost;
                    }
                    virtual-address [addresses];
                    virtual-link-local-address ipv6–address;
                    vrrp-inherit-from {
                      active-interface interface-name;
                      active-group group-number;
                    }
                  }
                }
              }
            }
```

Copyright © 2018, Juniper Networks, Inc.

Hierarchy Level [edit interfaces *interface-name*],
[edit logical-systems *logical-system-name* interfaces *interface-name*],
[edit interfaces interface-set *interface-set-name* interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options ***logical-unit-number***—Number of the logical unit.

Range: 0 through 1,073,741,823 for demux and PPPoE static interfaces. 0 through 16,385 for all other static interface types.

etree-ac-role (leaf | root)—To configure an interface as either leaf or root.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Logical Interface Properties*
- *Junos OS Services Interfaces Library for Routing Devices*

 Copyright © 2018, Juniper Networks, Inc.

unnumbered-address (Demux)

Syntax unnumbered-address *interface-name* <preferred-source-address *address*>;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet]

Release Information Statement introduced in Junos OS Release 8.2.
preferred-source-address option introduced in Junos OS Release 9.0.
IP demultiplexing interfaces supported in Junos OS Release 9.2.

Description For IP demultiplexing interfaces, enable the local address to be derived from the specified interface. Configuring an unnumbered interface enables IP processing on the interface without assigning an explicit IP address to the interface.

Options ***interface-name***—Name of the interface from which the local address is derived. The specified interface must have a logical unit number and a configured IP address, and must not be an unnumbered interface.

The **preferred-source-address** statement is explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring an Unnumbered Interface on page 218
- **address on page 424**
- *Junos System Basics Configuration Guide*

unnumbered-address (Dynamic Profiles)

Syntax unnumbered-address *interface-name* <preferred-source-address *address*>;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family*],
[edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family *family*]

Release Information Statement introduced in Junos OS Release 9.2.
Support for the **$junos-preferred-source-address** and **$junos-preferred-source-ipv6-address** predefined variables introduced in Junos OS Release 9.6.
Support for the **$junos-loopback-interface** predefined variable introduced in Junos OS Release 9.6.

Description For Ethernet interfaces, enable the local address to be derived from the specified interface. Configuring unnumbered Ethernet interfaces enables IP processing on the interface without assigning an explicit IP address to the interface. To configure unnumbered address dynamically, include the **$junos-loopback-interface-address** predefined variable.

You can configure unnumbered address support on Ethernet interfaces for IPv4 and IPv6 address families.

Options *interface-name*—Name of the interface from which the local address is derived. The specified interface must have a logical unit number, a configured IP address, and must not be an unnumbered interface. This value can be a specific interface name or the **$junos-loopback-interface** predefined variable.

When defining the **unnumbered-address** statement using a static interface, keep the following in mind:

- If you choose to include the **routing-instance** statement at the **[edit dynamic-profiles]** hierarchy level, that statement must be configured with a dynamic value by using the **$junos-routing-instance** predefined variable. In addition, whatever static unnumbered interface you specify must belong to that routing instance; otherwise, the profile instantiation fails.
- If you choose to not include the **routing-instance** statement at the **[edit dynamic-profiles]** hierarchy level, the unnumbered-address statement uses the default routing instance. The use of the default routing instance requires that the unnumbered interface be configured statically and that it reside in the default routing instance.

NOTE: When you specify a static logical interface for the unnumbered interface in a dynamic profile that includes the **$junos-routing-instance** predefined variable, you must not configure a preferred source address, whether with the **$junos-preferred-source-address** predefined variable, the **$junos-preferred-source-ipv6-address** predefined variable, or the

 Copyright © 2018, Juniper Networks, Inc.

preferred-source-address statement. Configuring the preferred source address in this circumstance causes a commit failure.

When defining the **unnumbered-address** statement using the **$junos-loopback-interface** predefined variable, keep the following in mind:

- To use the **$junos-loopback-interface** predefined variable, the dynamic profile must also contain the **routing-instance** statement configured with the **$junos-routing-instance** predefined variable at the [edit dynamic-profiles] hierarchy level.
- The applied loopback interface is based on the dynamically obtained routing instance of the subscriber.

address—(Optional) Secondary IP address of the donor interface. Configuring the preferred source address enables you to use an IP address other than the primary IP address on some of the unnumbered Ethernet interfaces in your network. This value can be a static IP address, the **$junos-preferred-source-address** predefined variable for the inet family, or the **$junos-preferred-source-ipv6-address** predefined variable for the inet6 family.

When defining the **preferred-source-address** value using a static IP address, keep the following in mind:

- The unnumbered interface must be statically configured.
- The IP address specified as the **preferred-source-address** must be configured in the specified unnumbered interface.

When defining the **preferred-source-address** value using the **$junos-preferred-source-address** or the **$junos-preferred-source-ipv6-address** predefined variables, keep the following in mind:

- You must configure the **unnumbered-address** statement using the **$junos-loopback-interface** predefined variable.
- You must configure the **routing-instance** statement using the **$junos-routing-instance** predefined variable at the **[edit dynamic-profiles]** hierarchy level.
- The preferred source address chosen is based on the dynamically applied loopback address which is in turn derived from the dynamically obtained routing instance of the subscriber. The configured loopback address with the closest network match to the user IP address is selected as the preferred source address.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Dynamic Profiles Overview*

unnumbered-address (Ethernet)

Syntax unnumbered-address *interface-name* <preferred-source-address *address*>;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family *family*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

Release Information Statement introduced in Junos OS Release 8.2.
preferred-source-address option introduced in Junos OS Release 9.0.

Description For Ethernet interfaces, enable the local address to be derived from the specified interface. Configuring an unnumbered Ethernet interface enables IP processing on the interface without assigning an explicit IP address to the interface.

Options ***interface-name***—Name of the interface from which the local address is derived. The specified interface must have a logical unit number and a configured IP address, and must not be an unnumbered interface.

The **preferred-source-address** statement is explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring an Unnumbered Interface on page 218
- address on page 424
- *Junos System Basics Configuration Guide*

 Copyright © 2018, Juniper Networks, Inc.

unnumbered-address (PPP)

Syntax unnumbered-address *interface-name* destination *address* destination-profile *profile-name*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet]

Release Information Statement introduced before Junos OS Release 7.4.

Description For interfaces with PPP encapsulation, enable the local address to be derived from the specified interface.

Options *interface-name*—Interface from which the local address is derived. The interface name must include a logical unit number and must have a configured address.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring IPCP Options for Interfaces with PPP Encapsulation on page 216
- *Junos OS Administration Library*

Copyright © 2018, Juniper Networks, Inc.

up-count

Syntax

up-count *cells*;

Hierarchy Level

[edit interfaces *interface-name* atm-options vpi *vpi-identifier* oam-liveness],
[edit interfaces *interface-name* unit *logical-unit-number* oam-liveness],
[edit interfaces *interface-name* unit *logical-unit-number* family *family* address *address* multipoint-destination *address* oam-liveness],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* oam-liveness],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family* address *address* multipoint-destination *address* oam-liveness]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

For ATM encapsulation only, configure Operation, Administration, and Maintenance (OAM) F5 loopback cell count thresholds. Not supported on ATM-over-SHDSL interfaces.

For ATM2 IQ PICs only, configure OAM F4 loopback cell count thresholds at the **[edit interfaces *interface-name* atm-options vpi *vpi-identifier*]** hierarchy level.

Options

cells—Minimum number of consecutive OAM F4 or F5 loopback cells received before a VC is declared up.

Range: 1 through 255

Default: 5 cells

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring the ATM OAM F5 Loopback Cell Threshold*

Copyright © 2018, Juniper Networks, Inc.

user-prefix

Syntax user-prefix *user-prefix-string*;

Hierarchy Level [edit interfaces *interface-name* auto-configure vlan-ranges authentication username-include],
[edit interfaces *interface-name* auto-configure stacked-vlan-ranges authentication username-include]

Release Information Statement introduced in Junos OS Release 10.0.

Description Specify the user prefix that is concatenated with the username during the subscriber authentication process.

Options ***user-prefix-string***—The user prefix string.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring VLAN Interface Username Information for AAA Authentication*

Copyright © 2018, Juniper Networks, Inc.

username-include

Syntax

```
username-include {
  circuit-id;
  circuit-type;
  delimiter delimiter-character;
  domain-name domain-name-string;
  interface-name;
  mac-address;
  option-18;
  option-37;
  option-82 <circuit-id> <remote-id>;
  radius-realm radius-realm-string;
  remote-id;
  user-prefix user-prefix-string;
}
```

Hierarchy Level

[edit interfaces *interface-name* auto-configure vlan-ranges authentication],
[edit interfaces *interface-name* auto-configure stacked-vlan-ranges authentication]

Release Information

Statement introduced in Junos OS Release 10.0.

Description

Configure the username that the router passes to the external AAA server. You must include at least one of the optional statements for the username to be valid. If you do not configure a username, the router accesses the local authentication service only and does not use external authentication services, such as RADIUS.

The username takes the format ***user-prefix mac-address circuit-type circuit-id remote-id option–82 interface-name domain-name radius-realm***. By default, each component is separated by a period (.), but you can specify a different delimiter with the **delimiter** statement.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring VLAN Interface Username Information for AAA Authentication*
- *Using DHCP Option 82 Suboptions in Authentication Usernames for Autosense VLANs*
- *Using DHCP Option 18 and Option 37 in Authentication Usernames for DHCPv6 Autosense VLANs*
- *Configuring a Username for Authentication of Out-of-Band Triggered Dynamic VLANs*

 Copyright © 2018, Juniper Networks, Inc.

vbr

Syntax vbr peak *rate* sustained *rate* burst *length*;

Hierarchy Level [edit interfaces *interface-name* atm-options vpi *vpi-identifier* shaping],
[edit interfaces *interface-name* unit *logical-unit-number* address *address* family *family* multipoint-destination *address* shaping],
[edit interfaces *interface-name* unit *logical-unit-number* shaping],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* address *address* family *family* multipoint-destination *address* shaping],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* shaping]

Release Information Statement introduced before Junos OS Release 7.4.

Description For ATM encapsulation only, define the variable bandwidth utilization in the traffic-shaping profile.

When you configure the variable bandwidth utilization, you must specify all three options (**burst**, **peak**, and **sustained**). You can specify the rate 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 the rate 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.

Default If the **vbr** statement is not specified, bandwidth utilization is unlimited.

Options **burst *length***—Burst length, in cells. If you set the length to 1, the peak traffic rate is used.

Range: 1 through 4000 cells

peak *rate*—Peak rate, in bits per second or cells per second.

Range: For ATM1 interfaces, 33 Kbps through 135.6 Mbps (ATM OC3); 33 Kbps through 276 Mbps (ATM OC12). For ATM2 IQ OC3 and OC12 interfaces, 33 Kbps through 542,526,792 bps. For ATM2 IQ OC48 interfaces, 33 Kbps through 2,170,107,168 bps. For ATM2 IQ DS3 and E3 interfaces, from 33 Kbps through the maximum rate, which depends on the ATM encapsulation and framing you configure.

sustained *rate*—Sustained rate, in bits per second or cells per second.

Range: For ATM1 interfaces, 33 Kbps through 135.6 Mbps (ATM OC3); 33 Kbps through 276 Mbps (ATM OC12). For ATM2 IQ OC3 and OC12 interfaces, 33 Kbps through 542,526,792 bps. For ATM2 IQ OC48 interfaces, 33 Kbps through 2,170,107,168 bps. For ATM2 IQ DS3 and E3 interfaces, from 33 Kbps through the maximum rate, which depends on the ATM encapsulation and framing you configure.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- *Configuring ATM CBR*
- *Applying Scheduler Maps to Logical ATM Interfaces*
- cbr on page 478
- rtvbr on page 944
- shaping on page 965

vc-cos-mode

Syntax vc-cos-mode (alternate | strict);

Hierarchy Level [edit interfaces *interface-name* atm-options scheduler-maps *map-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For ATM2 IQ interfaces only, specify packet-scheduling priority value for ATM2 IQ VC tunnels.

Options

alternate—VC CoS queue has high priority. The scheduling of the queues alternates between the high-priority queue and the remaining queues, so every other scheduled packet is from the high-priority queue.

strict—VC CoS queue has strictly high priority. A queue with strict high priority is always scheduled before the remaining queues. The remaining queues are scheduled in round-robin fashion.

Default: alternate

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *ATM2 IQ VC Tunnel CoS Components Overview*
- *Applying Scheduler Maps to ATM Interfaces*

 Copyright © 2018, Juniper Networks, Inc.

vci

Syntax vci *vpi-identifier*.*vci-identifier*;

Hierarchy Level [edit interfaces at-*fpc/pic/port* unit *logical-unit-number*],
[edit interfaces at-*fpc/pic/port* unit *logical-unit-number* family *family* address *address* multipoint-destination *address*],
[edit logical-systems *logical-system-name* interfaces at-*fpc/pic/port* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces at-*fpc/pic/port* unit *logical-unit-number* family *family* address *address* multipoint-destination *address*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro routers.

Description For ATM point-to-point logical interfaces only, configure the virtual circuit identifier (VCI) and virtual path identifier (VPI).

To configure a VPI for a point-to-multipoint interface, specify the VPI in the multipoint-destination statement.

VCIs 0 through 31 are reserved for specific ATM values designated by the ATM Forum.

Options ***vci-identifier***—ATM virtual circuit identifier. Unless you configure the interface to use promiscuous mode, this value cannot exceed the highest-numbered VC configured for the interface with the **maximum-vcs** option of the **vpi** statement.

Range: 0 through 4089 or 0 through 65,535 with promiscuous mode, with VCIs 0 through 31 reserved.

vpi-identifier—ATM virtual path identifier.

Range: 0 through 255

Default: 0

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring a Point-to-Point ATM1 or ATM2 IQ Connection*
- *Applying Scheduler Maps to Logical ATM Interfaces*

vci-range

Syntax vci-range start *start-vci* end *end-vci*;

Hierarchy Level [edit interfaces at-*fpc/pic/port* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces at-*fpc/pic/port* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.0.

Description Range of VCI values used in ATM-to-Ethernet interworking cross-connects. VCI 0 through 31 are reserved. VCI 0 through 31 should not be used.

Options ***start-vci***—Lowest number VCI in the range.

end-vci—Highest number VCI in the range.

Range: 0 through 255

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring ATM-to-Ethernet Interworking on page 282

virtual-switch

Syntax virtual-switch *name* bridge-domain *name* vlan-id [*vlan-ids*];

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* default-*x*]

Release Information Statement introduced in Junos OS Release 9.6.

Description Specify the routing-instance type as a virtual switch, under which bridge-domain MIPs must be enabled.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring MIP for Bridge Domains of a Virtual Switch*

 Copyright © 2018, Juniper Networks, Inc.

vlan-id (Logical Port in Bridge Domain)

Syntax vlan-id *number*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family bridge],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family bridge]

Release Information Statement introduced in Junos OS Release 9.2.
Statement introduced in Junos OS Release 15.1.

Description The VLAN ID configured on the logical port. Received packets with no VLAN tags are forwarded within the bridge domain with the matching VLAN ID.

Options number—The VLAN ID.

Range: 1 through 4095

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Access Mode on a Logical Interface*
- *Tunnel Services Overview*
- *Tunnel Interface Configuration on MX Series Routers Overview*

Copyright © 2018, Juniper Networks, Inc.

vlan-id (Outer VLAN ID)

Syntax vlan-id *outer-vlan-id*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.0.

Description The outer VLAN ID to be used in ATM-to-Ethernet interworking cross-connects. Outer VLAN IDs are converted to the ATM VPI. The outer VLAN ID must match the VPI value configured. The allowable VPI range is 0 to 255. Do not configure the outer VLAN ID to be greater than 255.

Options **outer-vlan-id**—Outer VLAN ID number.

Range: 0 through 4094

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring ATM-to-Ethernet Interworking on page 282

 Copyright © 2018, Juniper Networks, Inc.

vlan-id (VLAN ID to Be Bound to a Logical Interface)

Syntax vlan-id *number*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For Fast Ethernet, Gigabit Ethernet, and Aggregated Ethernet interfaces only, bind a 802.1Q VLAN tag ID to a logical interface.

Options *number*—A valid VLAN identifier.

Range: For aggregated Ethernet, 4-port, 8-port, and 12-port Fast Ethernet PICs, and for management and internal Ethernet interfaces, 1 through 1023.

For 48-port Fast Ethernet and Gigabit Ethernet PICs, 1 through 4094.

VLAN ID 0 is reserved for tagging the priority of frames.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Enabling VLAN Tagging*

Copyright © 2018, Juniper Networks, Inc.

native-vlan-id

Syntax

native-vlan-id *vlan-id*;

Hierarchy Level (QFX Series and EX4600)

For platforms without ELS:

[edit interfaces (QFX Series) *interface-name* unit 0 family ethernet-switching]

For platforms with ELS:

[edit interfaces (QFX Series) *interface-name*]

Hierarchy Level (ACX Series, EX Series, SRX Series, M Series, MX Series, and T Series)

[edit interfaces ge-*fpc/pic/port*],
[edit interfaces *interface-name*]

Hierarchy Level (SRX Series)

[edit interfaces*interface-name*]

Release Information

Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 9.5 for SRX Series.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.

Description

Configure the VLAN identifier to associate with untagged packets received on the physical interface of a trunk mode interface for the following:

- QFX Series and EX4600
- M Series routers with Gigabit Ethernet IQ PICs with SFP and Gigabit Ethernet IQ2 PICs with SFP configured for 802.1Q flexible VLAN tagging
- MX Series routers with Gigabit Ethernet DPCs and MICs, Tri-Rate Ethernet DPCs and MICs, and 10-Gigabit Ethernet DPCs and MICs and MPCs configured for 802.1Q flexible VLAN tagging
- T4000 routers with 100-Gigabit Ethernet Type 5 PIC with CFP
- EX Series switches with Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces

The logical interface on which untagged packets are received must be configured with the same VLAN ID as the native VLAN ID configured on the physical interface, otherwise the untagged packets are dropped. To configure the logical interface, include the **vlan-id** statement (matching the **native-vlan-id** statement on the physical interface) at the **[edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level**.

Copyright © 2018, Juniper Networks, Inc.

When the **native-vlan-id** statement is included with the **flexible-vlan-tagging** statement, untagged packets are accepted on the same mixed VLAN-tagged port and on the interfaces that are configured for Q-in-Q tunneling.

When the **native-vlan-id** statement is combined with the *interface-mode* statement, untagged packets are accepted and forwarded within the bridge domain or VLAN that is configured with the matching VLAN ID.

To configure the logical interface, include the **vlan-id** statement (matching the **native-vlan-id** statement on the physical interface) at the **[edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level**.

NOTE: **Starting in Junos OS Release 17.1R1, you can send untagged traffic without a native VLAN ID to the remote end of the network. To do this, remove the native VLAN ID from the untagged traffic configuration by setting the no-native-vlan-insert statement. If you do not configure this statement, the native VLAN ID is added to the untagged traffic.**

Default By default, the untagged packets are dropped. That is, if you do not configure the **native-vlan-id** option, the untagged packets are dropped.

Options *vlan-id*—Numeric identifier of the VLAN.

Range: 1 through 4094

number—VLAN ID number.

Range: (ACX Series routers, SRX Series devices and EX Series switches) 0 through 4094.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit Ethernet Interfaces (J-Web Procedure)*
- *Understanding Bridging and VLANs on Switches*
- *Enabling VLAN Tagging*
- *Configuring Access Mode on a Logical Interface*
- *Configuring the Native VLAN Identifier on Switches With ELS Support (CLI Procedure)*
- *Understanding Interfaces*
- *Understanding Q-in-Q Tunneling and VLAN Translation*
- *no-native-vlan-insert*
- *Sending Untagged Traffic Without VLAN ID to Remote End*
- *show ethernet-switching interfaces*
- *show vlans*
- flexible-vlan-tagging on page 618
- *Junos OS Network Interfaces Configuration Guide*

Copyright © 2018, Juniper Networks, Inc.

vlan-id-list (Ethernet VLAN Circuit)

Syntax vlan-id-list [*vlan-id vlan-id–vlan-id*];

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Binds a single-tag logical interface to a list of VLAN IDs. Configures a logical interface to receive and forward any tag frame whose VLAN ID tag matches the list of VLAN IDs you specify.

NOTE:
When you create a circuit cross-connect (CCC) using VLAN-bundled single-tag logical interfaces on Layer 2 VPN routing instances, the circuit automatically uses **ethernet** encapsulation. For Layer 2 VPN, you need to include the **encapsulation-type** statement and specify the value **ethernet** at either of the following hierarchy levels:

- **[edit routing-instances** ***routing-instance-name*** **protocols l2vpn]**
- **[edit logical-systems** ***logical-system-name*** **routing-instances** ***routing-instance-name*** **protocols l2vpn]**

For more information about the **encapsulation-type** configuration statement and the Layer 2 encapsulation types **ethernet** and **ethernet-vlan**, see the *Junos OS VPNs Library for Routing Devices*.

Options **[*vlan-id vlan-id–vlan-id*]**—A list of valid VLAN ID numbers. Specify the VLAN IDs individually by using a space to separate each ID, as an inclusive list by separating the starting VLAN ID and ending VLAN ID with a hyphen, or as a combination of both.

Range: 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

NOTE: Configuring **vlan-id-list** with the entire vlan-id range is an unnecessary waste of system resources and is not best practice. It should be used only when a subset of VLAN IDs (not the entire range) needs to be associated with a logical interface. If you specify the entire range (1-4094), it has the same result as not specifying a range; however, it consumes PFE resources such as VLAN lookup tables entries, and so on.

The following examples illustrate this further:

```
[edit interfaces interface-name]
vlan-tagging;
```

Copyright © 2018, Juniper Networks, Inc.

```
unit number {
    vlan-id-range 1-4094;
}

[edit interfaces interface-name]
unit 0;
```

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Binding VLAN IDs to Logical Interfaces*
- encapsulation (Logical Interface) on page 573
- encapsulation on page 577
- encapsulation-type (Layer 2 VPN routing instance), see the *Junos OS VPNs Library for Routing Devices*
- flexible-vlan-tagging on page 618
- vlan-tagging on page 1089
- vlan-tags (Dual-Tagged Logical Interface) on page 1091

 Copyright © 2018, Juniper Networks, Inc.

vlan-id-list (Interface in Bridge Domain)

Syntax vlan-id-list [*number number-number*];

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family bridge],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family bridge]

Release Information Statement introduced in Junos OS Release 9.2.
Statement introduced in Junos OS Release 15.1.

Description Configure a logical interface to forward packets and learn MAC addresses within each bridge domain configured with a VLAN ID that matches a VLAN ID specified in the list. VLAN IDs can be entered individually using a space to separate each ID, entered as an inclusive list separating the starting VLAN ID and ending VLAN ID with a hyphen, or a combination of both.

Options ***number number***—Individual VLAN IDs separated by a space.

number-number—Starting VLAN ID and ending VLAN ID in an inclusive range.

Range: 1 through 4095

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring a Logical Interface for Trunk Mode*
- *Configuring the VLAN ID List for a Trunk Interface*
- *Tunnel Services Overview*
- *Tunnel Interface Configuration on MX Series Routers Overview*

vlan-id-range

Syntax vlan-id-range *vlan-id–vlan-id*

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Bind a range of VLAN IDs to a logical interface.

Options **number**—The first number is the lowest VLAN ID in the range the second number is the highest VLAN ID in the range.

Range: 1 through 4094

NOTE: Configuring **vlan-id-range** with the entire vlan-id range is an unnecessary waste of system resources and is not best practice. It should be used only when a subset of VLAN IDs (not the entire range) needs to be associated with a logical interface. If you specify the entire range (1-4094), it has the same result as not specifying a range; however, it consumes PFE resources such as VLAN lookup tables entries, and so on.

The following examples illustrate this further:

```
[edit interfaces interface-name]
vlan-tagging;
unit number {
    vlan-id-range 1-4094;
}

[edit interfaces interface-name]
unit 0;
```

VLAN ID 0 is reserved for tagging the priority of frames.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Binding a Range of VLAN IDs to a Logical Interface*

 Copyright © 2018, Juniper Networks, Inc.

vlan-ranges

Syntax

```
vlan-ranges {
  access-profile profile-name;
  authentication {
    packet-types [packet-types];
    password password-string;
    username-include {
      circuit-type;
      circuit-id;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-name;
      mac-address;
      option-18;
      option-37;
      option-82 <circuit-id> <remote-id>;
      radius-realm radius-realm-string;
      remote-id;
      user-prefix user-prefix-string;
    }
  }
  dynamic-profile profile-name {
    accept (any | dhcp-v4 | inet);
    accept-out-of-band protocol;
    access-profilevlan-dynamic-profile-name;
    ranges (any | low-tag)–(any | high-tag);
  }
  override;
}
```

Hierarchy Level [edit interfaces *interface-name* auto-configure]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure multiple VLANs. Each VLAN is assigned a VLAN ID number from the range.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation

- *Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs*
- *Configuring Interfaces to Support Both Single and Stacked VLANs*

Copyright © 2018, Juniper Networks, Inc.

vlan-rewrite

Syntax vlan-rewrite translate (200 500 | 201 501)

Hierarchy Level [edit interfaces *interface-name* unit *number* family bridge interface-mode trunk]
[edit interfaces *interface-name* unit *number* family ethernet-switching interface-mode trunk]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Translates an incoming VLAN to a bridge-domain VLAN, corresponding counter translation at egress. Supports translation of VLAN 200 to VLAN 500 and VLAN 201 to VLAN 501. Other valid VLANs pass through without translation.

Options **translate 200 500**—Translates incoming packets with VLAN 200 to 500.

translate 201 501—Translates incoming packets with VLAN 201 to 501.

translate 202 502—Translates incoming packets with VLAN 202 to 502.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Rewriting a VLAN Tag and Adding a New Tag*

 Copyright © 2018, Juniper Networks, Inc.

vlan-rule (100-Gigabit Ethernet Type 4 PIC with CFP)

Syntax vlan-rule (high-low | odd-even);

Hierarchy Level [edit chassis fpc *slot* pic *slot* forwarding-mode vlan-steering]

Release Information Statement introduced in Junos OS Release 10.4.

Description Configure the interoperation mode of the 100-Gigabit Ethernet Type 4 PIC with CFP (PD-1CE-CFP-FPC4) when interoperating with 100 gigabit Ethernet interfaces from other vendors.

If no VLAN rule is configured, all tagged packets are distributed to PFE0.

Options **high-low**—VLAN IDs 1 through 2047 are distributed to PFE0 and VLAN IDs 2048 through 4096 are distributed to PFE1.

odd-even—Odd number VLAN IDs are distributed to PFE1 and even number VLAN IDs are distributed to PFE0.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP*
- *forwarding-mode (100-Gigabit Ethernet)*
- vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP) on page 1088

vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP)

Syntax

```
vlan-steering {
  vlan-rule (high-low | odd-even);
}
```

Hierarchy Level [edit chassis fpc *slot* pic *slot* forwarding-mode]

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure the 100-Gigabit Ethernet Type 4 PIC with CFP (PD-1CE-CFP-FPC4) to interoperate with 100 gigabit Ethernet interfaces from other vendors.

The other statement is explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP*
- *forwarding-mode (100-Gigabit Ethernet)*
- sa-multicast (100-Gigabit Ethernet) on page 946
- vlan-rule (100-Gigabit Ethernet Type 4 PIC with CFP) on page 1087

 Copyright © 2018, Juniper Networks, Inc.

vlan-tagging

Syntax vlan-tagging;

Syntax (QFX Series, NFX Series, and EX4600) vlan-tagging;

Syntax (SRX Series Interfaces) vlan-tagging native-vlan-id *vlan-id*;

Hierarchy Level [edit interfaces *interface-name*],
[edit logical-systems *logical-system-name* interfaces *interface-name*]

QFX Series, NFX Series, and EX4600 Interfaces [edit interfaces (QFX Series) *interface-name*]
[edit interfaces (QFX Series) interface-range *interface-range-name*]

SRX Series Interfaces [edit interfaces *interface*]

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 9.5.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.
Statement introduced in Junos OS Release 13.2 for PTX Series Routers.
Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series.

Description For Fast Ethernet and Gigabit Ethernet interfaces, aggregated Ethernet interfaces configured for VPLS, and pseudowire subscriber interfaces, enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface.

NOTE: For QFX Series configure VLAN identifier for untagged packets received on the physical interface of a trunk mode interface. Enable VLAN tagging. The platform receives and forwards single-tag frames with 802.1Q VLAN tags.

On EX Series switches except for EX4300 and EX9200 switches, the **vlan-tagging** and **family ethernet-switching** statements cannot be configured on the same interface. Interfaces on EX2200, EX3200, EX3300, EX4200, and EX4500 switches are set to **family ethernet-switching** by the default factory configuration. EX6200 and EX8200 switch interfaces do not have a default **family** setting.

Default VLAN tagging is disabled by default.

Copyright © 2018, Juniper Networks, Inc.

SRX Series [Warning: element unresolved in stylesheets: <title> (in <config-options>). This is probably a new element that is not yet supported in the stylesheets.]
SRX Series

native-vlan-id—Configures a VLAN identifier for untagged packets. Enter a number from 0 through 4094.

NOTE: The **native-vlan-id** can be configured only when either **flexible-vlan-tagging** mode or **interface-mode** trunk is configured.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *802.1Q VLANs Overview*
- *Configuring a Layer 3 Subinterface (CLI Procedure)*
- *Configuring Tagged Aggregated Ethernet Interfaces*
- *Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch*
- *vlan-id*
- *Configuring a Layer 3 Logical Interface*
- *Configuring VLAN Tagging*

 Copyright © 2018, Juniper Networks, Inc.

vlan-tags (Dual-Tagged Logical Interface)

Syntax vlan-tags inner-list [*vlan-id vlan-id–vlan-id*] outer <*tpid*.>*vlan-id*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.5.

Description (MX Series routers only) Binds a dual-tag logical interface to a list of VLAN IDs. Configures the logical interface to receive and forward any dual-tag frame whose inner VLAN ID tag matches the list of VLAN IDs you specify.

NOTE:
To create a circuit cross-connect (CCC) using VLAN-bundled dual-tag logical interfaces on Layer 2 VPN routing instances, you must include the **encapsulation-type** statement and specify the value **ethernet-vlan** at the one of the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* protocols l2vpn]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols l2vpn]**

For more information about the **encapsulation-type** configuration statement and the Layer 2 encapsulation types **ethernet** and **ethernet-vlan**, see the *Junos OS VPNs Library for Routing Devices*.

Options **inner-list [*vlan-id vlan-id vlan-id-vlan-id*]**—A list of valid VLAN ID numbers. Specify the VLAN IDs individually by using a space to separate each ID, as an inclusive list by separating the starting VLAN ID and ending VLAN ID with a hyphen, or as a combination of both.

Range: 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

outer <*tpid*.>*vlan-id*—An optional Tag Protocol ID (TPID) and a valid VLAN ID.

Range: For TPID, specify a hexadecimal value in the format **0x*nnnn***.

Range: For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

NOTE: Configuring **inner-list** with the entire vlan-id range is an unnecessary waste of system resources and is not best practice. It should be used only when a subset of VLAN IDs of inner tag (not the entire range) needs to be associated with a logical interface. If you specify the entire range (1 through

Copyright © 2018, Juniper Networks, Inc.

4094), it has the same result as not specifying a range; however, it consumes PFE resources such as VLAN lookup tables entries, and so on.

The following examples illustrate this further:

```
[edit interfaces interface-name]
vlan-tagging;
unit number {
    vlan-tags outer vid inner-list 1-4094;
}

[edit interfaces interface-name]
vlan-tagging;
unit number {
    vlan-id vid;
}
```

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Binding VLAN IDs to Logical Interfaces*
- encapsulation (Logical Interface) on page 573
- encapsulation on page 577
- encapsulation-type (Layer 2 VPN routing instance), see the *Junos OS VPNs Library for Routing Devices*.
- flexible-vlan-tagging on page 618
- vlan-id-list (Ethernet VLAN Circuit) on page 1081
- vlan-tagging on page 1089

 Copyright © 2018, Juniper Networks, Inc.

vlan-tags (Stacked VLAN Tags)

Syntax vlan-tags inner *tpid.vlan-id* inner-list *value* inner-range *vid1—vid2* outer *tpid.vlan-id*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

Description Bind TPIDs and 802.1Q VLAN tag IDs to a logical interface. TPID fields are used to identify the frame as an IEEE 802.1Q-tagged frame.

Options **inner** ***tpid.vlan-id***—A TPID and a valid VLAN identifier. TPID is a 16-bit field set to a value of 0x8100 in order to identify the frame as an IEEE 802.1Q-tagged frame.

Range: (most routers) For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames. For PTX Series, VLAN ID 0 is not supported.

inner-list ***value***— List or a set of VLAN identifiers.

NOTE: This is supported on MX Series routers with Trio-based FPCs.

inner-range ***tpid. vid1–vid2***—Specify a TPID and a range of VLAN IDs where vid1 is the start of the range and vid2 is the end of the range.

NOTE: On the network-to-network (NNI) or egress interfaces of provider edge (PE) routers, you cannot configure the **inner-range** ***tpid. vid1—vid2*** option with the **vlan-tags** statement for ISP-facing interfaces.

Range: For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

outer ***tpid.vlan-id***—A TPID and a valid VLAN identifier.

Range: (most routers) For VLAN ID, 1 through 511 for normal interfaces, and 512 through 4094 for VLAN CCC interfaces. VLAN ID 0 is reserved for tagging the priority of frames. For PTX Series, VLAN ID 0 is not supported.

NOTE: Configuring **inner-range** with the entire vlan-id range consumes system resources and is not a best practice. The **inner-range** must be used only when a subset of VLAN IDs of inner tag (not the entire range) needs to be associated with a logical interface. If you specify the entire range (1 through 4094), it

Copyright © 2018, Juniper Networks, Inc.

has the same result as not specifying a range; however, it consumes Packet Forwarding Engine resources such as VLAN lookup table entries, and so on.

The following examples illustrate this further:

```
[edit interfaces  interface-name]
stacked-vlan-tagging;
unit number {
    vlan-tags outer vid inner-range 1-4094;
}

[edit interfaces interface-name]
vlan-tagging;
unit number {
    vlan-id vid;
}
```

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Dual VLAN Tags*
- *Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers*
- stacked-vlan-tagging on page 987

vlan-tags-outer

Syntax

vlan-tags-outer *vlan-tag*;

Hierarchy Level

[edit interfaces interface-set *interface-set-name* interface *interface-name*]

Release Information

Statement introduced in Junos OS Release 8.5.

Description

The S-VLAN outer tag that belongs to a set of interfaces used to configure hierarchical CoS schedulers.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Class of Service Feature Guide for Routing Devices and EX9200 Switches*

Copyright © 2018, Juniper Networks, Inc.

vlan-vci-tagging

Syntax vlan-vci-tagging;

Hierarchy Level [edit interfaces *interface-name*],
[edit logical-systems *logical-system-name* interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 9.0.

Description Enable the ATM-to-Ethernet interworking cross-connect function on a Gigabit Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet interface.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring ATM-to-Ethernet Interworking on page 282

vpi (ATM CCC Cell-Relay Promiscuous Mode)

Syntax vpi *vpi-identifier*;

Hierarchy Level [edit interfaces at-*fpc/pic/port* atm-options promiscuous-mode]

Release Information Statement introduced before Junos OS Release 7.4.
Junos OS Release 12.2 for the ACX Series Universal Metro routers.

Description For ATM interfaces, allow all VCIs in this VPI to open in ATM CCC cell-relay mode.

When you include **vpi** statements at the **[edit interfaces *interface-name* atm-options promiscuous-mode]** hierarchy level, the specified VPIs open in promiscuous mode.

Options ***vpi-identifier***—ATM virtual path identifier. This is one of the VPIs that you define in the vci statement. (For a list of hierarchy levels at which you can include the **vci** statement, see vci.)

Range: 0 through 255

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring ATM Cell-Relay Promiscuous Mode*

Copyright © 2018, Juniper Networks, Inc.

vpi (Define Virtual Path)

Syntax

```
vpi vpi-identifier {
  maximum-vcs maximum-vcs;
  oam-liveness {
    up-count cells;
    down-count cells;
  }
  oam-period (disable | seconds);
  shaping {
    (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate burst
      length);
    queue-length number;
  }
}
```

Hierarchy Level [edit interfaces at-*fpc/pic/port* atm-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For ATM interfaces, configure the virtual path (VP).

NOTE: Certain options apply only to specific platforms.

Options *vpi-identifier*—ATM virtual path identifier. This is one of the VPIs that you define in the vci statement. (For a list of hierarchy levels at which you can include the **vci** statement, see vci.)

Range: 0 through 255

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Maximum Number of ATM1 VCs on a VP*

 Copyright © 2018, Juniper Networks, Inc.

vpi (Logical Interface and Interworking)

Syntax vpi *virtual-path-identifier*;

Hierarchy Level [edit interfaces at-*fpc/pic/port* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces at-*fpc/pic/port* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro routers.

Description VPI used in an ATM-to-Ethernet interworking cross-connect.

Options **virtual-path-identifier**—VPI to be used.

Range: 0 through 255

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring ATM-to-Ethernet Interworking on page 282
- *Configuring ATM Cell-Relay Promiscuous Mode*

vtmapping

Syntax

vtmapping (itu-t | klm);

Hierarchy Level

[edit chassis fpc *number* pic *number*],
[edit interfaces *interface-name* sonet-options]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

For the Channelized STM1 IQ PIC or Channelized STM1 PIC, configure virtual tributary mapping.

For the Channelized STM1 PIC, you configure virtual tributary mapping at the **[edit chassis fpc *number* pic *number*]** hierarchy level.

NOTE: The **vtmapping** statement is not supported for **cau4** interfaces on the Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (H).

Options

itu-t—International Telephony Union standard.

klm—KLM standard.

Default: klm

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Virtual Tributary Mapping of Channelized STM1 Interfaces*
- *Configuring the Junos OS to Support Channelized STM1 Interface Virtual Tributary Mapping*

 Copyright © 2018, Juniper Networks, Inc.

warning

Syntax

```
warning low-light-warning {
  (link-down | syslog);
}
```

Hierarchy Level [edit interfaces *interface-name* optics-options]

Release Information Statement introduced in Junos OS Release 10.0.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Description Specifies the action to take if the receiving optics signal is below the optics low-light warning threshold.

Options **link-down**—Drop the 10-Gigabit Ethernet link and marks link as down.

syslog—Write the optics information to the system log.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring 10-Gigabit Ethernet Link Down Notification for Optics Options Alarm or Warning*
- optics-options on page 839
- *100-Gigabit Ethernet OTN Options Configuration Overview*

watch-list

Syntax

```
watch-list {
  [ routes ];
}
```

Hierarchy Level [edit interfaces dl*n* unit *logical-unit-number* dialer-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description On J Series Services Routers with ISDN interfaces, configure an ISDN list of routes to watch. Used only for dialer watch.

Options ***routes***—IP prefix of a route. Specify one or more. The primary interface is considered up if there is at least one valid route for any of the addresses in the watch list to an interface other than the backup interface.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Junos OS Interfaces and Routing Configuration Guide*

Copyright © 2018, Juniper Networks, Inc.

wavelength

Syntax wavelength *nm*;

Hierarchy Level [edit interfaces *interface-name* optics-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Description For 10-Gigabit or 100-Gigabit Ethernet DWDM interfaces only, configure full C-band ITU-Grid tunable optics.

Options *nm*—Wavelength value. It can be one of the following:

NOTE: All values are displayed. However, if you configure a value that is not supported by the device, an error message is displayed and the device is not tuned to the specified wavelength.

- **1528.38**—1528.38 nanometers (nm), corresponds to a 50-GHz grid
- **1528.77**—1528.77 nm, corresponds to 50-GHz and 100-GHz grids
- **1529.16**—1529.16 nm, corresponds to a 50-GHz grid
- **1529.55**—1529.55 nm, corresponds to 50-GHz and 100-GHz grids
- **1529.94**—1529.94 nm, corresponds to a 50-GHz grid
- **1530.33**—1530.33 nm, corresponds to 50-GHz and 100-GHz grids
- **1530.72**—1530.72 nm, corresponds to a 50-GHz grid
- **1531.12**—1531.12 nm, corresponds to 50-GHz and 100-GHz grids
- **1531.51**—1531.51 nm, corresponds to a 50-GHz grid
- **1531.90**—1531.90 nm, corresponds to 50-GHz and 100-GHz grids
- **1532.29**—1532.29 nm, corresponds to a 50-GHz grid
- **1532.68**—1532.68 nm, corresponds to 50-GHz and 100-GHz grids
- **1533.07**—1533.07 nm, corresponds to a 50-GHz grid
- **1533.47**—1533.47 nm, corresponds to 50-GHz and 100-GHz grids
- **1533.86**—1533.86 nm, corresponds to a 50-GHz grid
- **1534.25**—1534.25 nm, corresponds to 50-GHz and 100-GHz grids
- **1534.64**—1534.64 nm, corresponds to a 50-GHz grid

Copyright © 2018, Juniper Networks, Inc.

- **1535.04**—1535.04 nm, corresponds to 50-GHz and 100-GHz grids
- **1535.43**—1535.43 nm, corresponds to a 50-GHz grid
- **1535.82**—1535.82 nm, corresponds to 50-GHz and 100-GHz grids
- **1536.22**—1536.22 nm, corresponds to a 50-GHz grid
- **1536.61**—1536.61 nm, corresponds to 50-GHz and 100-GHz grids
- **1537.00**—1537.00 nm, corresponds to a 50-GHz grid
- **1537.40**—1537.40 nm, corresponds to 50-GHz and 100-GHz grids
- **1537.79**—1537.79 nm, corresponds to a 50-GHz grid
- **1538.19**—1538.19 nm, corresponds to 50-GHz and 100-GHz grids
- **1538.58**—1538.58 nm, corresponds to a 50-GHz grid
- **1538.98**—1538.98 nm, corresponds to 50-GHz and 100-GHz grids
- **1539.37**—1539.37 nm, corresponds to a 50-GHz grid
- **1539.77**—1539.77 nm, corresponds to 50-GHz and 100-GHz grids
- **1540.16**—1540.16 nm, corresponds to a 50-GHz grid
- **1540.56**—1540.56 nm, corresponds to 50-GHz and 100-GHz grids
- **1540.95**—1540.95 nm, corresponds to a 50-GHz grid
- **1541.35**—1541.35 nm, corresponds to 50-GHz and 100-GHz grids
- **1541.75**—1541.75 nm, corresponds to a 50-GHz grid
- **1542.14**—1542.14 nm, corresponds to 50-GHz and 100-GHz grids
- **1542.54**—1542.54 nm, corresponds to a 50-GHz grid
- **1542.94**—1542.94 nm, corresponds to 50-GHz and 100-GHz grids
- **1543.33**—1543.33 nm, corresponds to a 50-GHz grid
- **1543.73**—1543.73 nm, corresponds to 50-GHz and 100-GHz grids
- **1544.13**—1544.13 nm, corresponds to a 50-GHz grid
- **1544.53**—1544.53 nm, corresponds to 50-GHz and 100-GHz grids
- **1544.92**—1544.92 nm, corresponds to a 50-GHz grid
- **1545.32**—1545.32 nm, corresponds to 50-GHz and 100-GHz grids
- **1545.72**—1545.72 nm, corresponds to a 50-GHz grid
- **1546.12**—1546.12 nm, corresponds to 50-GHz and 100-GHz grids
- **1546.52**—1546.52 nm, corresponds to a 50-GHz grid
- **1546.92**—1546.92 nm, corresponds to 50-GHz and 100-GHz grids
- **1547.32**—1547.32 nm, corresponds to a 50-GHz grid
- **1547.72**—1547.72 nm, corresponds to 50-GHz and 100-GHz grids

Copyright © 2018, Juniper Networks, Inc.

- **1548.11**—1548.11 nm, corresponds to a 50-GHz grid
- **1548.51**—1548.51 nm, corresponds to 50-GHz and 100-GHz grids
- **1548.91**—1548.91 nm, corresponds to a 50-GHz grid
- **1549.32**—1549.32 nm, corresponds to 50-GHz and 100-GHz grids
- **1549.72**—1549.72 nm, corresponds to a 50-GHz grid
- **1550.12**—1550.12 nm, corresponds to 50-GHz and 100-GHz grids
- **1550.52**—1550.52 nm, corresponds to a 50-GHz grid
- **1550.92**—1550.92 nm, corresponds to 50-GHz and 100-GHz grids
- **1551.32**—1551.32 nm, corresponds to a 50-GHz grid
- **1551.72**—1551.72 nm, corresponds to 50-GHz and 100-GHz grids
- **1552.12**—1552.12 nm, corresponds to a 50-GHz grid
- **1552.52**—1552.52 nm, corresponds to 50-GHz and 100-GHz grids
- **1552.93**—1552.93 nm, corresponds to a 50-GHz grid
- **1553.33**—1554.33 nm, corresponds to 50-GHz and 100-GHz grids
- **1553.73**—1554.73 nm, corresponds to a 50-GHz grid
- **1554.13**—1554.13 nm, corresponds to 50-GHz and 100-GHz grids
- **1554.54**—1554.54 nm, corresponds to a 50-GHz grid
- **1554.94**—1554.94 nm, corresponds to 50-GHz and 100-GHz grids
- **1555.34**—1555.34 nm, corresponds to a 50-GHz grid
- **1555.75**—1555.75 nm, corresponds to 50-GHz and 100-GHz grids
- **1556.15**—1556.15 nm, corresponds to a 50-GHz grid
- **1556.55**—1556.55 nm, corresponds to 50-GHz and 100-GHz grids
- **1556.96**—1556.96 nm, corresponds to a 50-GHz grid
- **1557.36**—1557.36 nm, corresponds to 50-GHz and 100-GHz grids
- **1557.77**—1557.77 nm, corresponds to a 50-GHz grid
- **1558.17**—1558.17 nm, corresponds to 50-GHz and 100-GHz grids
- **1558.58**—1558.58 nm, corresponds to a 50-GHz grid
- **1558.98**—1558.98 nm, corresponds to 50-GHz and 100-GHz grids
- **1559.39**—1559.39 nm, corresponds to a 50-GHz grid
- **1559.79**—1559.79 nm, corresponds to 50-GHz and 100-GHz grids
- **1560.20**—1560.20 nm, corresponds to a 50-GHz grid
- **1560.61**—1560.61 nm, corresponds to 50-GHz and 100-GHz grids
- **1561.01**—1561.01 nm, corresponds to a 50-GHz grid

Copyright © 2018, Juniper Networks, Inc.

- **1561.42**—1561.42 nm, corresponds to 50-GHz and 100-GHz grids
- **1561.83**—1561.83 nm, corresponds to a 50-GHz grid
- **1562.23**—1562.23 nm, corresponds to 50-GHz and 100-GHz grids
- **1562.64**—1562.64 nm, corresponds to a 50-GHz grid
- **1563.05**—1563.05 nm, corresponds to 50-GHz and 100-GHz grids
- **1563.45**—1563.45 nm, corresponds to a 50-GHz grid
- **1563.86**—1563.86 nm, corresponds to 50-GHz and 100-GHz grids
- **1564.27**—1564.27 nm, corresponds to a 50-GHz grid
- **1564.68**—1564.68 nm, corresponds to 50-GHz and 100-GHz grids
- **1565.09**—1565.09 nm, corresponds to a 50-GHz grid
- **1565.50**—1565.50 nm, corresponds to 50-GHz and 100-GHz grids
- **1565.90**—1565.90 nm, corresponds to a 50-GHz grid
- **1566.31**—1566.31 nm, corresponds to 50-GHz and 100-GHz grids
- **1566.72**—1566.72 nm, corresponds to a 50-GHz grid
- **1567.13**—1567.13 nm, corresponds to 50-GHz and 100-GHz grids
- **1567.54**—1567.54 nm, corresponds to a 50-GHz grid
- **1567.95**—1567.95 nm, corresponds to 50-GHz and 100-GHz grids
- **1568.36**—1568.36 nm, corresponds to a 50-GHz grid
- **1568.77**—1568.77 nm, corresponds to 50-GHz and 100-GHz grids

Default: 1550.12—1550.12 nm, corresponds to 50-GHz and 100-GHz grids

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Ethernet DWDM Interface Wavelength Overview*
- *Configuring the 10-Gigabit or 100-Gigabit Ethernet DWDM Interface Wavelength*
- *show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port)*

 Copyright © 2018, Juniper Networks, Inc.

west-interface

Syntax

```
west-interface {
  node-id mac-address;
  control-channel channel-name {
    vlan number;
    interface name interface-name
  }
  interface-none
  ring-protection-link-end;
  virtual-control-channel {
    west-interface name;
    east-interface name;
  }
}
```

Hierarchy Level

[edit protocols protection-group ethernet-ring *ring-name*]

Release Information

Statement introduced in Junos OS Release 9.5.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

Description

Define one of the two interface ports for Ethernet ring protection, the other being defined by the **east-interface** statement at the same hierarchy level. The interface must use the control channel's logical interface name. The control channel is a dedicated VLAN channel for the ring port.

NOTE: Always configure this port second, after configuring the **east-interface** statement.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Ethernet Ring Protection Switching Overview*
- *Ethernet Ring Protection Using Ring Instances for Load Balancing*
- east-interface on page 570
- ethernet-ring on page 595
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

Copyright © 2018, Juniper Networks, Inc.

working-circuit

Syntax working-circuit *group-name*;

Hierarchy Level [edit interfaces *interface-name* sonet-options aps]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the working router in an APS circuit pair.

Options *group-name*—Circuit's group name.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Basic Automatic Protect Switching*
- protect-circuit on page 906

yellow-differential-delay

Syntax yellow-differential-delay *milliseconds*;

Hierarchy Level [edit interfaces *interface-name* mlfr-uni-nni-bundle-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description For link services and voices interfaces only, configure the yellow differential delay among bundle links to give warning when a link has a differential delay that exceeds the configured threshold.

Options *milliseconds*—Yellow differential delay threshold.

Range: 1 through 2000 milliseconds

Default: 6 milliseconds

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*
- action-red-differential-delay on page 421
- remote on page 926

 Copyright © 2018, Juniper Networks, Inc.

z0-increment

Syntax (z0-increment | no-z0-increment);

Hierarchy Level [edit interfaces *interface-name* sonet-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure an incremental STM ID rather than a static one.

Default no-Z0-increment

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring an Incrementing STM ID to Interoperate with Older Equipment in SDH Mode*
- sonet-options on page 973

Copyright © 2018, Juniper Networks, Inc.

CHAPTER 14

Interface Operational Commands

- show interfaces lsi (Label-Switched Interface)
- show interfaces media
- show interfaces terse

Common Output Fields Description

This chapter explains the content of the output fields, which appear in the output of most **show interfaces** commands.

Damping Field

For the physical interface, the Damping field shows the setting of the following damping parameters:

- **half-life**—Decay half-life. The number of seconds after which the accumulated interface penalty counter is reduced by half if the interface remains stable.
- **max-suppress**—Maximum hold-down time. The maximum number of seconds that an interface can be suppressed irrespective of how unstable the interface has been.
- **reuse**—Reuse threshold. When the accumulated interface penalty counter falls below this number, the interface is no longer suppressed.
- **suppress**—Cutoff (suppression) threshold. When the accumulated interface penalty counter exceeds this number, the interface is suppressed.
- **state**—Interface damping state. If damping is enabled on an interface, it is suppressed during interface flaps that match the configured damping parameters.

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:

```
                                           Packets                        Bytes
Destination class                  (packet-per-second)          (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:

 Copyright © 2018, Juniper Networks, Inc.

```
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

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.
- **Trunk**—Interface is a trunk.
- **Trunk, Inter-Switch-Link**—Interface is a trunk, and InterSwitch Link protocol (ISL) is configured on the trunk port of the primary VLAN in order to connect the routers composing the PVLAN to each other.

Device Flags Field

The **Device flags** field provides information about the physical device and displays one or more of the following values:

- **ASIC Error**—Device is down because of ASIC wedging and due to which PFE is disabled.
- **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.

 Copyright © 2018, Juniper Networks, Inc.

- **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.
- **OAM-On-SVLAN**—(MX Series routers with MPC/MIC interfaces only) Interface is configured to propagate the Ethernet OAM state of a static, single-tagged service VLAN (S-VLAN) on a Gigabit Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet interface to a dynamic or static double-tagged customer VLAN (C-VLAN) that has the same S-VLAN (outer) tag as the S-VLAN.
- **Point-To-Point**—Interface is point-to-point.

Copyright © 2018, Juniper Networks, Inc.

- **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.

 Copyright © 2018, Juniper Networks, Inc.

- **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

Copyright © 2018, Juniper Networks, Inc.

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.

 Copyright © 2018, Juniper Networks, Inc.

- **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:

```
                     Packets               Bytes
Source class   (packet-per-second)    (bits-per-second)

        gold        1928095            161959980
           (            889)     (       597762)
      bronze              0                    0
           (              0)     (             0)
      silver              0                    0
           (              0)     (             0)
```

Improvements to Interface Transmit Statistics Reporting

The offered load on an interface can be defined as the amount of data the interface is capable of transmitting during a given time period. The actual traffic that goes out of the interface is the transmitted load. However, when outgoing interfaces are oversubscribed, there could be traffic drops in the schedulers attached to the outgoing interfaces. Hence, the offered load is not always the same as the actual transmitted load because the offered load calculation does not take into account possible packet drop or traffic loss.

On MX Series routers, the logical interface-level statistics show the offered load, which is often different from the actual transmitted load. To address this limitation, Junos OS introduces a new configuration option in Release 11.4 R3 and later. The new configuration option, **interface-transmit-statistics**, at the **[edit interface *interface-name*]** hierarchy level, enables you to configure Junos OS to accurately capture and report the transmitted load on interfaces.

When the **interface-transmit-statistics** statement is included at the **[edit interface *interface-name*]** hierarchy level, the following operational mode commands report the actual transmitted load:

- **show interface *interface-name* <detail | extensive>**

- **monitor interface** ***interface-name***
- **show snmp mib get** ***objectID.ifIndex***

The **show interface** ***interface-name*** command also shows whether the **interface-transmit-statistics** configuration is enabled or disabled on the interface.

Related Documentation

- interface-transmit-statistics on page 689
- show interfaces on page 1119

 Copyright © 2018, Juniper Networks, Inc.

show interfaces

List of Syntax

Syntax (Gigabit Ethernet) on page 1119
Syntax (10 Gigabit Ethernet) on page 1119
Syntax (SRX Series Devices) on page 1119

Syntax (Gigabit Ethernet)

show interfaces ge-*fpc/pic/port*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Syntax (10 Gigabit Ethernet)

show interfaces xe-*fpc/pic/port*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Syntax (SRX Series Devices)

show interfaces (
  <*interface-name*>
  <brief | detail | extensive | terse>
  <controller *interface-name*>|
  <descriptions *interface-name*>|
  <destination-class (all | *destination-class-name logical-interface-name*)>|
  <diagnostics optics *interface-name*>|
  <far-end-interval *interface-fpc/pic/port*>|
  <filters *interface-name*>|
  <flow-statistics *interface-name*>|
  <interval *interface-name*>|
  <load-balancing (detail | *interface-name*)>|
  <mac-database mac-address *mac-address*>|
  <mc-ae id *identifier* unit *number* revertive-info>|
  <media *interface-name*>|
  <policers *interface-name*>|
  <queue both-ingress-egress egress forwarding-class *forwarding-class* ingress l2-statistics>|
  <redundancy (detail | *interface-name*)>|
  <routing brief detail summary *interface-name*>|
  <routing-instance (all | *instance-name*)>|
  <snmp-index *snmp-index*>|
  <source-class (all | *destination-class-name logical-interface-name*)>|
  <statistics *interface-name*>|
  <switch-port *switch-port number*>|
  <transport pm (all | optics | otn) (all | current | currentday | interval | previousday) (all | *interface-name*)>|
  <zone *interface-name*>
)

Release Information

Command introduced before Junos OS Release 7.4 for Gigabit interfaces.
Command introduced in Junos OS Release 8.0 for 10 Gigabit interfaces.
Command modified in Junos OS Release 9.5 for SRX Series devices.

Copyright © 2018, Juniper Networks, Inc.

Command introduced in Junos OS Release 18.1 for Gigabit interfaces.

Description Display status information about the specified Gigabit Ethernet interface.

(M320, M120, MX Series, and T Series routers only) Display status information about the specified 10-Gigabit Ethernet interface.

Display the IPv6 interface traffic statistics about the specified Gigabit Ethernet interface for MX series routers. The input and output bytes (bps) and packets (pps) rates are not displayed for IFD and local traffic.

Display status information and statistics about interfaces on SRX Series appliance running Junos OS.

NOTE: On SRX Series appliances, on configuring identical IPs on a single interface, you will not see a warning message; instead, you will see a syslog message.

Options For Gigabit interfaces:

ge-*fpc/pic/port*—Display standard information about the specified Gigabit Ethernet interface.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

For 10 Gigabit interfaces:

xe-*fpc/pic/port*—Display standard information about the specified 10-Gigabit Ethernet interface.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

For SRX interfaces:

Copyright © 2018, Juniper Networks, Inc.

- **interface-name**—(Optional) Display standard information about the specified interface. Following is a list of typical interface names. Replace pim with the PIM slot and port with the port number.
  - **at-** ***pim*****/0/*****port***—ATM-over-ADSL or ATM-over-SHDSL interface.
  - **ce1-*****pim*****/0/** ***port***—Channelized E1 interface.
  - **cl-0/0/8**—3G wireless modem interface for SRX320 devices.
  - **ct1-*****pim*****/0/*****port***—Channelized T1 interface.
  - **dl0**—Dialer Interface for initiating ISDN and USB modem connections.
  - **e1-*****pim*****/0/*****port***—E1 interface.
  - **e3-*****pim*****/0/*****port***—E3 interface.
  - **fe-*****pim*****/0/*****port***—Fast Ethernet interface.
  - **ge-*****pim*****/0/*****port***—Gigabit Ethernet interface.
  - **se-*****pim*****/0/*****port***—Serial interface.
  - **t1-*****pim*****/0/*****port***—T1 (also called DS1) interface.
  - **t3-*****pim*****/0/*****port***—T3 (also called DS3) interface.
  - **wx-*****slot*****/0/0**—WAN acceleration interface, for the WXC Integrated Services Module (ISM 200).
  - **interface-name**—(Optional) Display standard information about the specified interface. Following is a list of typical interface names. Replace pim with the PIM slot and port with the port number.
    - **at-** ***pim*****/0/*****port***—ATM-over-ADSL or ATM-over-SHDSL interface.
    - **ce1-*****pim*****/0/** ***port***—Channelized E1 interface.
    - **cl-0/0/8**—3G wireless modem interface for SRX320 devices.
    - **ct1-*****pim*****/0/*****port***—Channelized T1 interface.
    - **dl0**—Dialer Interface for initiating ISDN and USB modem connections.
    - **e1-*****pim*****/0/*****port***—E1 interface.
    - **e3-*****pim*****/0/*****port***—E3 interface.
    - **fe-*****pim*****/0/*****port***—Fast Ethernet interface.
    - **ge-*****pim*****/0/*****port***—Gigabit Ethernet interface.
    - **se-*****pim*****/0/*****port***—Serial interface.
    - **t1-*****pim*****/0/*****port***—T1 (also called DS1) interface.

- **t3-*pim/0/port***—T3 (also called DS3) interface.
- **wx-*slot*/0/0**—WAN acceleration interface, for the WXC Integrated Services Module (ISM 200).

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

- *Understanding Layer 2 Interfaces on Security Devices*
- *Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*
- *Verifying and Managing Configurations for Dynamic VLANs Based on Access-Line Identifiers*

List of Sample Output

 Copyright © 2018, Juniper Networks, Inc.

Output Fields Table 56 on page 1123 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 57 on page 1152.

Table 56: 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 1110.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none

Table 56: 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**
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Loopback	Loopback status: **Enabled** or **Disabled**. If loopback is enabled, type of loopback: **Local** or **Remote**.	All levels
Source filtering	Source filtering status: **Enabled** or **Disabled**.	All levels
LAN-PHY mode	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
WAN-PHY mode	10-Gigabit Ethernet interface operating in Wide Area Network Physical Layer Device (WAN PHY) mode. WAN PHY allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and other devices intended for SONET/SDH.	All levels
Unidirectional	Unidirectional link mode status for 10-Gigabit Ethernet interface: **Enabled** or **Disabled** for parent interface; **Rx-only** or **Tx-only** for child interfaces.	All levels
Flow control	Flow control status: **Enabled** or **Disabled**.	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: **Enabled** or **Disabled**.	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: • **Online**—Autonegotiation is manually configured as online. • **Offline**—Autonegotiation is manually configured as offline.	All levels
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under "Common Output Fields Description" on page 1110.	All levels
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under "Common Output Fields Description" on page 1110.	All levels
Link flags	Information about the link. Possible values are described in the "Links Flags" section under "Common Output Fields Description" on page 1110.	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

 Copyright © 2018, Juniper Networks, Inc.

Table 56: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
Schedulers	(Gigabit Ethernet intelligent queuing 2 [IQ2] interfaces only) Number of CoS schedulers configured.	**extensive**
Hold-times	Current interface hold-time up and hold-time down, in milliseconds (ms).	**detail extensive**
Current address	Configured MAC address.	**detail extensive** none
Hardware address	Hardware MAC address.	**detail extensive** none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour:minute:second:timezone*** **(*hour:minute:second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).**	**detail extensive** none
Input Rate	Input rate in bits per second (bps) and packets per second (pps). The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None
Output Rate	Output rate in bps and pps. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Egress account overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for egress traffic.	**detail extensive**
Ingress account overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for ingress traffic.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Output bytes**—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface. Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes. For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 31 under the show interfaces command.	**detail extensive**

Table 56: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • **L3 incompletes**—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the **ignore-l3-incompletes** statement. • **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **FIFO errors**—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **Resource errors**—Sum of transmit drops.	**extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 56: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the **Drops** field does not always use the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs. • **Collisions**—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number must always be 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 must 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. NOTE: In DPCs that are not of the enhanced type, such as DPC 40x 1GE R, DPCE 20x 1GE + 2x 10GE R, or DPCE 40x 1GE R, you might notice a discrepancy in the output of the **show interfaces** command because incoming packets might be counted in the Egress queues section of the output. This problem occurs on non-enhanced DPCs because the egress queue statistics are polled from IMQ (Inbound Message Queuing) block of the I-chip. The IMQ block does not differentiate between ingress and egress WAN traffic; as a result, the combined statistics are displayed in the egress queue counters on the Routing Engine. In a simple VPLS scenario, if there is no MAC entry in DMAC table (by sending unidirectional traffic), traffic is flooded and the input traffic is accounted in IMQ. For bidirectional traffic (MAC entry in DMAC table), if the outgoing interface is on the same I-chip then both ingress and egress statistics are counted in a combined way. If the outgoing interface is on a different I-chip or FPC, then only egress statistics are accounted in IMQ. This behavior is expected with non-enhanced DPCs	**detail extensive**

Table 56: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
Queue counters (Egress)	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism. NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the **Dropped packets** field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.	**detail extensive**
Ingress queues	Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.	**extensive**
Queue counters (Ingress)	CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**extensive**
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value **None** or **Link**. • **None**—There are no active defects or alarms. • **Link**—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	**detail extensive** none
Interface transmit statistics	(On MX Series devices) Status of the **interface-transmit-statistics** configuration: Enabled or Disabled. • **Enabled**—When the **interface-transmit-statistics** statement is included in the configuration. If this is configured, the interface statistics show the actual transmitted load on the interface. • **Disabled**—When the **interface-transmit-statistics** statement is not included in the configuration. If this is not configured, the interface statistics show the offered load on the interface.	**detail extensive**
OTN FEC statistics	The forward error correction (FEC) counters provide the following statistics: • **Corrected Errors**—Count of corrected errors in the last second. • **Corrected Error Ratio**—Corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits.	**detail extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 56: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
PCS statistics	(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device. • **Bit errors**—Number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode. • **Errored blocks**—Number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode.	**detail extensive**

Link Degrade	Shows the link degrade status of the physical link and the estimated bit error rates (BERs). This field is available only for the PICs supporting the physical link monitoring feature.
	• **Link Monitoring**—Indicates if physical link degrade monitoring is enabled on the interface.
	  ◦ **Enable**—Indicates that link degrade monitoring has been enabled (using the **link-degrade-monitor** statement) on the interface.
	  ◦ **Disable**—Indicates that link degrade monitoring has not been enabled on the interface. If link degrade monitoring has not been enabled, the output does not show any related information, such as BER values and thresholds.
	• **Link Degrade Set Threshold**—The BER threshold value at which the link is considered degraded and a corrective action is triggered.
	• **Link Degrade Clear Threshold**—The BER threshold value at which the degraded link is considered recovered and the corrective action applied to the interface is reverted.
	• **Estimated BER**—The estimated bit error rate.
	• **Link-degrade event**—Shows link degrade event information.
	  ◦ **Seconds**—Time (in seconds) elapsed after a link degrade event occurred.
	  ◦ **Count**—The number of link degrade events recorded.
	  ◦ **State**—Shows the link degrade status (example: **Defect Active**).
		detail extensive

Table 56: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output

MAC statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC subsystem, including the following:
	
	• **Total octets** and **total packets**—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see Table 31 under the show interfaces command.
	• **Unicast packets, Broadcast packets,** and **Multicast packets**—Number of unicast, broadcast, and multicast packets.
	• **CRC/Align errors**—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).
	• **FIFO error**—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning.
	• **MAC control frames**—Number of MAC control frames.
	• **MAC pause frames**—Number of MAC control frames with **pause** operational code.
	• **Oversized frames**—There are two possible conditions regarding the number of oversized frames:
	  ◦ Packet length exceeds 1518 octets, or
	  ◦ Packet length exceeds MRU
	• **Jabber frames**—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.
	• **Fragment frames**—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.
	• **VLAN tagged frames**—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not.
	
	NOTE: The 20-port Gigabit Ethernet MIC (MIC-3D-20GE-SFP) does not have hardware counters for VLAN frames. Therefore, the **VLAN tagged frames** field displays 0 when the **show interfaces** command is executed on a 20-port Gigabit Ethernet MIC. In other words, the number of VLAN tagged frames cannot be determined for the 20-port Gigabit Ethernet MIC.
	
	• **Code violations**—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error."
		extensive

OTN Received Overhead Bytes	APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08	**extensive**
OTN Transmitted Overhead Bytes	APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 56: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
Filter statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet may enter the system or be rejected. • **Input packet count**—Number of packets received from the MAC hardware that the filter processed. • **Input packet rejects**—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • **Input DA rejects**—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the router from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local router (which the router is rejecting). • **Input SA rejects**—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field must 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 must 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 must be 0.	**extensive**
PMA PHY	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information: • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. Any state other than **OK** indicates a problem. Subfields are: • **PHY Lock**—Phase-locked loop • **PHY Light**—Loss of optical signal	**extensive**

Table 56: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
WIS section	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information: • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. Any state other than **OK** indicates a problem. Subfields are: • **BIP-B1**—Bit interleaved parity for SONET section overhead • **SEF**—Severely errored framing • **LOL**—Loss of light • **LOF**—Loss of frame • **ES-S**—Errored seconds (section) • **SES-S**—Severely errored seconds (section) • **SEFS-S**—Severely errored framing seconds (section)	**extensive**
WIS line	(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information: • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. Any state other than **OK** indicates a problem. Subfields are: • **BIP-B2**—Bit interleaved parity for SONET line overhead • **REI-L**—Remote error indication (near-end line) • **RDI-L**—Remote defect indication (near-end line) • **AIS-L**—Alarm indication signal (near-end line) • **BERR-SF**—Bit error rate fault (signal failure) • **BERR-SD**—Bit error rate defect (signal degradation) • **ES-L**—Errored seconds (near-end line) • **SES-L**—Severely errored seconds (near-end line) • **UAS-L**—Unavailable seconds (near-end line) • **ES-LFE**—Errored seconds (far-end line) • **SES-LFE**—Severely errored seconds (far-end line) • **UAS-LFE**—Unavailable seconds (far-end line)	**extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 56: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
WIS path	(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information: • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. Any state other than **OK** indicates a problem. Subfields are: • **BIP-B3**—Bit interleaved parity for SONET section overhead • **REI-P**—Remote error indication • **LOP-P**—Loss of pointer (path) • **AIS-P**—Path alarm indication signal • **RDI-P**—Path remote defect indication • **UNEQ-P**—Path unequipped • **PLM-P**—Path payload (signal) label mismatch • **ES-P**—Errored seconds (near-end STS path) • **SES-P**—Severely errored seconds (near-end STS path) • **UAS-P**—Unavailable seconds (near-end STS path) • **SES-PFE**—Severely errored seconds (far-end STS path) • **UAS-PFE**—Unavailable seconds (far-end STS path)	**extensive**

Table 56: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output

Autonegotiation information	Information about link autonegotiation.
	• **Negotiation status:**
	  ◦ **Incomplete**—Ethernet interface has the speed or link mode configured.
	  ◦ **No autonegotiation**—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
	  ◦ **Complete**—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• **Link partner status**—**OK** when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• **Link partner**—Information from the remote Ethernet device:
	  ◦ **Link mode**—Depending on the capability of the link partner, either **Full-duplex** or **Half-duplex**.
	  ◦ **Flow control**—Types of flow control supported by the link partner. For Gigabit Ethernet interfaces, types are **Symmetric** (link partner supports **PAUSE** on receive and transmit), **Asymmetric** (link partner supports **PAUSE** on transmit), **Symmetric/Asymmetric** (link partner supports **PAUSE** on receive and transmit or only **PAUSE** on transmit), and **None** (link partner does not support flow control).
	  ◦ **Remote fault**—Remote fault information from the link partner—**Failure** indicates a receive link error. **OK** indicates that the link partner is receiving. **Negotiation error** indicates a negotiation error. **Offline** indicates that the link partner is going offline.
	• **Local resolution**—Information from the local Ethernet device:
	  ◦ **Flow control**—Types of flow control supported by the local device. For Gigabit Ethernet interfaces, advertised capabilities are **Symmetric/Asymmetric** (local device supports **PAUSE** on receive and transmit or only **PAUSE** on receive) and **None** (local device does not support flow control). Depending on the result of the negotiation with the link partner, local resolution flow control type will display **Symmetric** (local device supports **PAUSE** on receive and transmit), **Asymmetric** (local device supports **PAUSE** on receive), and **None** (local device does not support flow control).
	  ◦ **Remote fault**—Remote fault information. **Link OK** (no error detected on receive), **Offline** (local interface is offline), and **Link Failure** (link error detected on receive).
		extensive

Received path trace, Transmitted path trace	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the router at the other end of the fiber. The transmitted path trace value is the message that this router transmits.	**extensive**
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number.	**extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 56: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface. • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP interface index number for the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under "Common Output Fields Description" on page 1110.	All levels

Table 56: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
VLAN-Tag	Rewrite profile applied to incoming or outgoing frames on the outer (**Out**) VLAN tag or for both the outer and inner (**In**) VLAN tags. • **push**—An outer VLAN tag is pushed in front of the existing VLAN tag. • **pop**—The outer VLAN tag of the incoming frame is removed. • **swap**—The outer VLAN tag of the incoming frame is overwritten with the user-specified VLAN tag information. • **push**—An outer VLAN tag is pushed in front of the existing VLAN tag. • **push-push**—Two VLAN tags are pushed in from the incoming frame. • **swap-push**—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • **swap-swap**—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user-specified VLAN tag value. • **pop-swap**—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • **pop-pop**—Both the outer and inner VLAN tags of the incoming frame are removed.	**brief detail extensive** none
Demux	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: • Source Family Inet • Destination Family Inet	**detail extensive** none
Encapsulation	Encapsulation on the logical interface.	All levels
ACI VLAN	Information displayed for agent circuit identifier (ACI) interface set configured with the **agent-circuit-id** autoconfiguration stanza. **Dynamic Profile**—Name of the dynamic profile that defines the ACI interface set. If configured, the ACI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ACI information. NOTE: The ACI VLAN field is replaced with the Line Identity field when an ALI interface set is configured with the **line-identity** autoconfiguration stanza.	**brief detail extensive** none
Line Identity	Information displayed for access-line-identifier (ALI) interface sets configured with the **line-identity** autoconfiguration stanza. • **Dynamic Profile**—Name of the dynamic profile that defines the ALI interface set. • Trusted option used to create the ALI interface set: **Circuit-id**, **Remote-id**, or **Accept-no-ids**. More than one option can be configured. If configured, the ALI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ALI information. NOTE: The Line Identity field is replaced with the ACI VLAN field when an ACI interface set is configured with the **agent-circuit-id** autoconfiguration stanza.	**detail**

 Copyright © 2018, Juniper Networks, Inc.

Table 56: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under “Common Output Fields Description” on page 1110.	**detail extensive** none
MTU	Maximum transmission unit size on the logical interface.	**detail extensive** none
Neighbor Discovery Protocol (NDP)Queue Statistics	NDP statistics for protocol **inet6** under logical interface statistics. • **Max nh cache**—Maximum interface neighbor discovery nexthop cache size. • **New hold nh limit**—Maximum number of new unresolved nexthops. • **Curr nh cnt**—Current number of resolved nexthops in the NDP queue. • **Curr new hold cnt**—Current number of unresolved nexthops in the NDP queue. • **NH drop cnt**—Number of NDP requests not serviced.	All levels
Dynamic Profile	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	Name of the service name table for the interface configured with a PPPoE family.	**detail extensive** none
Max Sessions	Maximum number of PPPoE logical interfaces that can be activated on the underlying interface.	**detail extensive** none
Duplicate Protection	State of PPPoE duplicate protection: **On** or **Off**. When duplicate protection is configured for the underlying interface, a dynamic PPPoE logical interface cannot be activated when an existing active logical interface is present for the same PPPoE client.	**detail extensive** none
Direct Connect	State of the configuration to ignore DSL Forum VSAs: **On** or **Off**. When configured, the router ignores any of these VSAs received from a directly connected CPE device on the interface.	**detail extensive** none
AC Name	Name of the access concentrator.	**detail extensive** none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	**detail extensive** none
Traffic statistics	Number and rate of bytes and packets received and transmitted on the specified interface set. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Input packets, Output packets**—Number of packets received and transmitted on the interface set.	**detail extensive**
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	**extensive**
Local statistics	Number and rate of bytes and packets destined to the router.	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 56: 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 1110.	**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**
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 1110.	**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 1110.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none

Copyright © 2018, Juniper Networks, Inc.

Table 56: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
Broadcast	Broadcast address of the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

The following table describes the output fields for the **show interfaces** (10–Gigabit Ethernet) command.

Field Name	Field Description	Level of Output
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 1110.	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

Copyright © 2018, Juniper Networks, Inc.

Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: • **Online**—Autonegotiation is manually configured as online. • **Offline**—Autonegotiation is manually configured as offline.	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under “Common Output Fields Description” on page 1110.	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 1110.	All levels
Link flags	Information about the link. Possible values are described in the “Links Flags” section under “Common Output Fields Description” on page 1110.	All levels
Wavelength	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels
Frequency	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels
CoS queues	Number of CoS queues configured.	**detail extensive** none
Schedulers	(Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces only) Number of CoS schedulers configured.	**extensive**
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive** none
Hardware address	Hardware MAC address.	**detail extensive** none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour:minute:second:timezone*** **(*hour:minute:second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).**	**detail extensive** none
Input Rate	Input rate in bits per second (bps) and packets per second (pps). The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None specified
Output Rate	Output rate in bps and pps. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Egress account overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for egress traffic.	**detail extensive**
Ingress account overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for ingress traffic.	**detail extensive**

 Copyright © 2018, Juniper Networks, Inc.

Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Output bytes**—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • **L3 incompletes**—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the **ignore-l3-incompletes** statement. • **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **FIFO errors**—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **Resource errors**—Sum of transmit drops.	**extensive**

Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Collisions**—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • **Aged packets**—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • **FIFO errors**—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**
Egress queues	Total number of egress queues supported on the specified interface. NOTE: In DPCs that are not of the enhanced type, such as DPC 40x 1GE R, DPCE 20x 1GE + 2x 10GE R, or DPCE 40x 1GE R, you might notice a discrepancy in the output of the **show interfaces** command because incoming packets might be counted in the Egress queues section of the output. This problem occurs on non-enhanced DPCs because the egress queue statistics are polled from IMQ (Inbound Message Queuing) block of the I-chip. The IMQ block does not differentiate between ingress and egress WAN traffic; as a result, the combined statistics are displayed in the egress queue counters on the Routing Engine. In a simple VPLS scenorio, if there is no MAC entry in DMAC table (by sending unidirectional traffic), traffic is flooded and the input traffic is accounted in IMQ. For bidirectional traffic (MAC entry in DMAC table), if the outgoing interface is on the same I-chip then both ingress and egress statistics are counted in a combined way. If the outgoing interface is on a different I-chip or FPC, then only egress statistics are accounted in IMQ. This behavior is expected with non-enhanced DPCs	**detail extensive**
Queue counters (Egress)	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**
Ingress queues	Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.	**extensive**

 Copyright © 2018, Juniper Networks, Inc.

Queue counters (Ingress)	CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**extensive**
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the routing device configuration, an alarm can ring the red or yellow alarm bell on the routing device, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value **None** or **Link**. • **None**—There are no active defects or alarms. • **Link**—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	**detail extensive** none
OTN alarms	Active OTN alarms identified on the interface.	**detail extensive**
OTN defects	OTN defects received on the interface.	**detail extensive**
OTN FEC Mode	The FECmode configured on the interface. • **efec**—Enhanced forward error correction (EFEC) is configured to defect and correct bit errors. • **gfec**—G.709 Forward error correction (GFEC) mode is configured to detect and correct bit errors. • **none**—FEC mode is not configured.	**detail extensive**
OTN Rate	OTN mode. • **fixed-stuff-bytes**—Fixed stuff bytes 11.0957 Gbps. • **no-fixed-stuff-bytes**—No fixed stuff bytes 11.0491 Gbps. • **pass-through**—Enable OTN passthrough mode. • **no-pass-through**—Do not enable OTN passthrough mode.	**detail extensive**
OTN Line Loopback	Status of the line loopback, if configured for the DWDM OTN PIC. Its value can be: **enabled** or **disabled**.	**detail extensive**
OTN FEC statistics	The forward error correction (FEC) counters for the DWDM OTN PIC. • **Corrected Errors**—The count of corrected errors in the last second. • **Corrected Error Ratio**—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits.	**detail extensive**
OTN FEC alarms	OTN FEC excessive or degraded error alarms triggered on the interface. • **FEC Degrade**—OTU FEC Degrade defect. • **FEC Excessive**—OTU FEC Excessive Error defect.	**detail extensive**

OTN OC	OTN OC defects triggered on the interface. • **LOS**—OC Loss of Signal defect. • **LOF**—OC Loss of Frame defect. • **LOM**—OC Loss of Multiframe defect. • **Wavelength Lock**—OC Wavelength Lock defect.	**detail extensive**
OTN OTU	OTN OTU defects detected on the interface • **AIS**—OTN AIS alarm. • **BDI**—OTN OTU BDI alarm. • **IAE**—OTN OTU IAE alarm. • **TTIM**—OTN OTU TTIM alarm. • **SF**—OTN ODU bit error rate fault alarm. • **SD**—OTN ODU bit error rate defect alarm. • **TCA-ES**—OTN ODU ES threshold alarm. • **TCA-SES**—OTN ODU SES threshold alarm. • **TCA-UAS**—OTN ODU UAS threshold alarm. • **TCA-BBE**—OTN ODU BBE threshold alarm. • **BIP**—OTN ODU BIP threshold alarm. • **BBE**—OTN OTU BBE threshold alarm. • **ES**—OTN OTU ES threshold alarm. • **SES**—OTN OTU SES threshold alarm. • **UAS**—OTN OTU UAS threshold alarm.	**detail extensive**
Received DAPI	Destination Access Port Interface (DAPI) from which the packets were received.	**detail extensive**
Received SAPI	Source Access Port Interface (SAPI) from which the packets were received.	**detail extensive**
Transmitted DAPI	Destination Access Port Interface (DAPI) to which the packets were transmitted.	**detail extensive**
Transmitted SAPI	Source Access Port Interface (SAPI) to which the packets were transmitted.	**detail extensive**
PCS statistics	(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device. • **Bit errors**—The number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode. • **Errored blocks**—The number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode.	**detail extensive**

 Copyright © 2018, Juniper Networks, Inc.

MAC statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC subsystem, including the following: • **Total octets** and **total packets**—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. • **Unicast packets, Broadcast packets,** and **Multicast packets**—Number of unicast, broadcast, and multicast packets. • **CRC/Align errors**—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • **FIFO error**—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • **MAC control frames**—Number of MAC control frames. • **MAC pause frames**—Number of MAC control frames with **pause** operational code. • **Oversized frames**—Number of frames that exceed 1518 octets. • **Jabber frames**—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • **Fragment frames**—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • **VLAN tagged frames**—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. • **Code violations**—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”	extensive
OTN Received Overhead Bytes	APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08	extensive
OTN Transmitted Overhead Bytes	APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08	extensive

Copyright © 2018, Juniper Networks, Inc.

Filter statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected. • **Input packet count**—Number of packets received from the MAC hardware that the filter processed. • **Input packet rejects**—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • **Input DA rejects**—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the routing device from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local routing device (which the routing device is rejecting). • **Input SA rejects**—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • **Output packet count**—Number of packets that the filter has given to the MAC hardware. • **Output packet pad count**—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • **Output packet error count**—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment. • **CAM destination filters, CAM source filters**—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields should be 0.	extensive
PMA PHY	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information: • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. Any state other than **OK** indicates a problem.	extensive

Copyright © 2018, Juniper Networks, Inc.

WIS section	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information: • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. Any state other than **OK** indicates a problem. Subfields are: • **BIP-B1**—Bit interleaved parity for SONET section overhead • **SEF**—Severely errored framing • **LOL**—Loss of light • **LOF**—Loss of frame • **ES-S**—Errored seconds (section) • **SES-S**—Severely errored seconds (section) • **SEFS-S**—Severely errored framing seconds (section)	**extensive**
WIS line	(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **BIP-B2**—Bit interleaved parity for SONET line overhead • **REI-L**—Remote error indication (near-end line) • **RDI-L**—Remote defect indication (near-end line) • **AIS-L**—Alarm indication signal (near-end line) • **BERR-SF**—Bit error rate fault (signal failure) • **BERR-SD**—Bit error rate defect (signal degradation) • **ES-L**—Errored seconds (near-end line) • **SES-L**—Severely errored seconds (near-end line) • **UAS-L**—Unavailable seconds (near-end line) • **ES-LFE**—Errored seconds (far-end line) • **SES-LFE**—Severely errored seconds (far-end line) • **UAS-LFE**—Unavailable seconds (far-end line)	**extensive**

WIS path	(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. Any state other than **OK** indicates a problem. Subfields are: • **BIP-B3**—Bit interleaved parity for SONET section overhead • **REI-P**—Remote error indication • **LOP-P**—Loss of pointer (path) • **AIS-P**—Path alarm indication signal • **RDI-P**—Path remote defect indication • **UNEQ-P**—Path unequipped • **PLM-P**—Path payload label mismatch • **ES-P**—Errored seconds (near-end STS path) • **SES-P**—Severely errored seconds (near-end STS path) • **UAS-P**—Unavailable seconds (near-end STS path) • **SES-PFE**—Severely errored seconds (far-end STS path) • **UAS-PFE**—Unavailable seconds (far-end STS path)	**extensive**

Autonegotiation information	Information about link autonegotiation.
	• **Negotiation status:**
	  ◦ **Incomplete**—Ethernet interface has the speed or link mode configured.
	  ◦ **No autonegotiation**—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
	  ◦ **Complete**—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• **Link partner status**—**OK** when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• **Link partner:**
	  ◦ **Link mode**—Depending on the capability of the attached Ethernet device, either **Full-duplex** or **Half-duplex**.
	  ◦ **Flow control**—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is **None**. For Gigabit Ethernet interfaces, types are **Symmetric** (link partner supports **PAUSE** on receive and transmit), **Asymmetric** (link partner supports **PAUSE** on transmit), and **Symmetric/Asymmetric** (link partner supports both **PAUSE** on receive and transmit or only **PAUSE** receive).
	  ◦ **Remote fault**—Remote fault information from the link partner—**Failure** indicates a receive link error. **OK** indicates that the link partner is receiving. **Negotiation error** indicates a negotiation error. **Offline** indicates that the link partner is going offline.
	• **Local resolution**—Information from the link partner:
	  ◦ **Flow control**—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are **Symmetric** (link partner supports **PAUSE** on receive and transmit), **Asymmetric** (link partner supports **PAUSE** on transmit), and **Symmetric/Asymmetric** (link partner supports both **PAUSE** on receive and transmit or only **PAUSE** receive).
	  ◦ **Remote fault**—Remote fault information. **Link OK** (no error detected on receive), **Offline** (local interface is offline), and **Link Failure** (link error detected on receive).
		extensive

 Copyright © 2018, Juniper Networks, Inc.

Received path trace, Transmitted path trace	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the routing device at the other end of the fiber. The transmitted path trace value is the message that this routing device transmits.	**extensive**
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number.	**extensive**
CoS information	Information about the CoS queue for the physical interface. • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP interface index number for the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under “Common Output Fields Description” on page 1110.	All levels

VLAN-Tag	Rewrite profile applied to incoming or outgoing frames on the outer (**Out**) VLAN tag or for both the outer and inner (**In**) VLAN tags. • **push**—An outer VLAN tag is pushed in front of the existing VLAN tag. • **pop**—The outer VLAN tag of the incoming frame is removed. • **swap**—The outer VLAN tag of the incoming frame is overwritten with the user specified VLAN tag information. • **push**—An outer VLAN tag is pushed in front of the existing VLAN tag. • **push-push**—Two VLAN tags are pushed in from the incoming frame. • **swap-push**—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • **swap-swap**—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user specified VLAN tag value. • **pop-swap**—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • **pop-pop**—Both the outer and inner VLAN tags of the incoming frame are removed.	**brief detail extensive** none
Demux:	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: • Source Family Inet • Destination Family Inet	**detail extensive** none
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under “Common Output Fields Description” on page 1110.	**detail extensive** none
MTU	Maximum transmission unit size on the logical interface.	**detail extensive** none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	**detail extensive** none
Traffic statistics	Number and rate of bytes and packets received and transmitted on the specified interface set. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Input packets, Output packets**—Number of packets received and transmitted on the interface set.	**detail extensive**
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	**extensive**
Local statistics	Number and rate of bytes and packets destined to the routing device.	**extensive**

 Copyright © 2018, Juniper Networks, Inc.

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 1110.	**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 1110.	**detail extensive** none
protocol-family	Protocol family configured on the logical interface. If the protocol is **inet**, the IP address of the interface is also displayed.	**brief**
Flags	Information about address flag (possible values are described in the “Addresses Flags” section under “Common Output Fields Description” on page 1110.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interlace.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Copyright © 2018, Juniper Networks, Inc.

For Gigabit Ethernet IQ PICs, traffic and MAC statistics output varies. The following table 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). 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 57: Gigabit and 10 Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type

Interface Type	Sample Command	Byte and Octet Counts Include	Comments
Inbound physical interface	**show interfaces ge-0/3/0 extensive**	Traffic statistics: Input bytes: 496 bytes per packet, representing the Layer 2 packet MAC statistics: Received octets: 500 bytes per packet, representing the Layer 2 packet + 4 bytes	The additional 4 bytes are for the CRC.
Inbound logical interface	**show interfaces ge-0/3/0.50 extensive**	Traffic statistics: Input bytes: 478 bytes per packet, representing the Layer 3 packet	
Outbound physical interface	**show interfaces ge-0/0/0 extensive**	Traffic statistics: Input bytes: 490 bytes per packet, representing the Layer 3 packet + 12 bytes MAC statistics: Received octets: 478 bytes per packet, representing the Layer 3 packet	For input bytes, the additional 12 bytes 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	

Table 58 on page 1153 lists the output fields for the **show interfaces** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2018, Juniper Networks, Inc.

Table 58: show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface.	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
Link-level type	Encapsulation being used on the physical interface.	All levels
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
MTU	Maximum transmission unit size on the physical interface.	All levels
Link mode	Link mode: Full-duplex or Half-duplex.	
Speed	Speed at which the interface is running.	All levels
BPDU error	Bridge protocol data unit (BPDU) error: Detected or None	
Loopback	Loopback status: **Enabled** or **Disabled**. If loopback is enabled, type of loopback: **Local** or **Remote**.	All levels
Source filtering	Source filtering status: **Enabled** or **Disabled**.	All levels
Flow control	Flow control status: **Enabled** or **Disabled**.	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: **Enabled** or **Disabled**.	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: • **Online**—Autonegotiation is manually configured as online. • **Offline**—Autonegotiation is manually configured as offline.	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Link flags	Information about the physical link.	All levels
CoS queues	Number of CoS queues configured.	**detail extensive** none
Current address	Configured MAC address.	**detail extensive** none

Copyright © 2018, Juniper Networks, Inc.

Table 58: show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
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
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. These fields can contain the value **None** or **Link**. • **None**—There are no active defects or alarms. • **Link**—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	**detail extensive** none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 58: show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface. • **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** . • **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**
Output errors	Output errors on the interface. • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), 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; therefore, for Gigabit Ethernet PICs, this number must 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 must 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 interfaces. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 58: show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Ingress queues	Total number of ingress queues supported on the specified interface.	**extensive**
Queue counters and queue number	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**

MAC statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC subsystem, including the following:
	• **Total octets** and **total packets**—Total number of octets and packets.
	• **Unicast packets, Broadcast packets,** and **Multicast packets**—Number of unicast, broadcast, and multicast packets.
	• **CRC/Align errors**—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).
	• **FIFO error**—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning.
	• **MAC control frames**—Number of MAC control frames.
	• **MAC pause frames**—Number of MAC control frames with **pause** operational code.
	• **Oversized frames**—There are two possible conditions regarding the number of oversized frames:
	  ◦ Packet length exceeds 1518 octets, or
	  ◦ Packet length exceeds MRU
	• **Jabber frames**—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.
	• **Fragment frames**—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.
	• **VLAN tagged frames**—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not.
	• **Code violations**—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”
		extensive

 Copyright © 2018, Juniper Networks, Inc.

Table 58: show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Filter statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected. • **Input packet count**—Number of packets received from the MAC hardware that the filter processed. • **Input packet rejects**—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • **Input DA rejects**—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the device from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local device (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 must be 0.	**extensive**

Autonegotiation information	Information about link autonegotiation.
	• **Negotiation status:**
	  ◦ **Incomplete**—Ethernet interface has the speed or link mode configured.
	  ◦ **No autonegotiation**—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
	  ◦ **Complete**—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
		extensive

Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number.	**extensive**

Table 58: show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface. • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**
Interface transmit statistics	Status of the **interface-transmit-statistics** configuration: Enabled or Disabled.	**detail extensive**
Queue counters (Egress)	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**
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.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Traffic statistics	Number and rate of bytes and packets received and transmitted on the specified interface set. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Input packets, Output packets**—Number of packets received and transmitted on the interface set.	**detail extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 58: show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Local statistics	Number and rate of bytes and packets destined to the device.	**extensive**
Transit statistics	Number and rate of bytes and packets transiting the switch. NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the **Output bytes** and **Output packets** interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.	**extensive**
Security	Security zones that interface belongs to.	**extensive**
Flow Input statistics	Statistics on packets received by flow module.	**extensive**
Flow Output statistics	Statistics on packets sent by flow module.	**extensive**
Flow error statistics (Packets dropped due to)	Statistics on errors in the flow module.	**extensive**
Protocol	Protocol family.	**detail extensive** none
MTU	Maximum transmission unit size on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Route table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive** none
Flags	Information about protocol family flags. .	**detail extensive**
Addresses, Flags	Information about the address flags..	**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**

Copyright © 2018, Juniper Networks, Inc.

Sample Output Gigabit Ethernet

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: 00:00:5e:00:53:7c, Hardware address: 00:00:5e:00:53:7c
  Last flapped   : 2006-08-10 17:25:10 PDT (00:01:08 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Ingress rate at Packet Forwarding Engine      : 0 bps (0 pps)
  Ingress drop rate at Packet Forwarding Engine : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None

  Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69)
    Flags: SNMP-Traps 0x4000
    VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push
    0x8100.512 0x8100.513)
    Encapsulation: VLAN-CCC
    Egress account overhead: 100
    Ingress account overhead: 90
    Input packets : 0
    Output packets: 0
    Protocol ccc, MTU: 1522
      Flags: Is-Primary
```

show interfaces (Gigabit Ethernet on MX Series Routers)

```
user@host> show interfaces ge-2/2/2
Physical interface: ge-2/2/2, Enabled, Physical link is Up
  Interface index: 156, SNMP ifIndex: 188
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, MAC-REWRITE Error: None,
 Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Schedulers     : 0
  Current address: 00:00:5e:00:53:c0, Hardware address: 00:00:5e:00:53:76
  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: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
    Input packets : 10232
    Output packets: 10294
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
```

 Copyright © 2018, Juniper Networks, Inc.

```
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 203.0.113/24, Local: 203.0.113.1, Broadcast: 203.0.113.255
   Protocol inet6, MTU: 1500
    Max nh cache: 4, New hold nh limit: 100000, Curr nh cnt: 4, Curr new hold
cnt: 4, NH drop cnt: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 2001:db8:/32, Local: 2001:db8::5
      Addresses, Flags: Is-Preferred
        Destination: 2001:db8:1::/32, Local: 2001:db8:223:9cff:fe9f:3e78
    Protocol multiservice, MTU: Unlimited
      Flags: Is-Primary
```

show interfaces (link degrade status)

```
user@host> show interfaces et-3/0/0
Physical interface: et-3/0/0, Enabled, Physical link is Down
  Interface index: 157, SNMP ifIndex: 537
  Link-level type: Ethernet, MTU: 1514, MRU: 0, Speed: 100Gbps, BPDU 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
  Current address: 54:e0:32:23:9d:38, Hardware address: 54:e0:32:23:9d:38
  Last flapped   : 2014-06-18 02:36:38 PDT (02:50:50 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0
  Link Degrade* :
  Link Monitoring                   :  Enable
  Link Degrade Set Threshold:       :  1E-7
  Link Degrade Clear Threshold:     :  1E-12
  Estimated BER                     :  1E-7
  Link-degrade event                :  Seconds      Count     State
                                           782          1      Defect Active
```

show interfaces extensive (Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration)

```
user@host> show interfaces ge-2/1/2 extensive | match "output|interface"
Physical interface: ge-2/1/2, Enabled, Physical link is Up
  Interface index: 151, SNMP ifIndex: 530, Generation: 154
  Interface flags: SNMP-Traps Internal: 0x4000
   Output bytes  :         240614363944            772721536 bps
   Output packets:           3538446506              1420444 pps
    Direction : Output
  Interface transmit statistics: Enabled

  Logical interface ge-2/1/2.0 (Index 331) (SNMP ifIndex 955) (Generation 146)
     Output bytes  :         195560312716            522726272 bps
     Output packets:           4251311146              1420451 pps

user@host> show interfaces ge-5/2/0.0 statistics detail
Logical interface ge-5/2/0.0 (Index 71) (SNMP ifIndex 573) (Generation 135)
    Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
```

```
    Egress account overhead: 100
    Ingress account overhead: 90
    Traffic statistics:
     Input  bytes  :          271524
     Output bytes  :        37769598
     Input  packets:            3664
     Output packets:          885790
     IPv6 transit statistics:
      Input  bytes  :              0
      Output bytes  :       16681118
      Input  packets:              0
      Output packets:         362633
    Local statistics:
     Input  bytes  :          271524
     Output bytes  :          308560
     Input  packets:            3664
     Output packets:            3659
    Transit statistics:
     Input  bytes  :               0                  0 bps
     Output bytes  :        37461038                  0 bps
     Input  packets:               0                  0 pps
     Output packets:          882131                  0 pps
     IPv6 transit statistics:
      Input  bytes  :              0     0 bps
      Output bytes  :       16681118     0 bps
      Input  packets:              0     0 pps
      Output packets:         362633     0 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
```


Copyright © 2018, Juniper Networks, Inc.

```
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:7c, Hardware address: 00:00:5e:00:53:7c
  Last flapped   : 2006-08-09 17:17:00 PDT (01:31:33 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                    0 bps
   Output bytes  :                   0                    0 bps
   Input  packets:                   0                    0 pps
   Output packets:                   0                    0 pps
  Ingress traffic statistics at Packet Forwarding Engine:
   Input  bytes  :                   0                    0 bps
   Input  packets:                   0                    0 pps
   Drop   bytes  :                   0                    0 bps
   Drop   packets:                   0                    0 pps
  Ingress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                   0                    0                    0

    1 expedited-fo                  0                    0                    0

    2 assured-forw                  0                    0                    0

    3 network-cont                  0                    0                    0

  Egress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                   0                    0                    0

    1 expedited-fo                  0                    0                    0

    2 assured-forw                  0                    0                    0

    3 network-cont                  0                    0                    0

  Active alarms  : None
  Active defects : None

  Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69) (Generation 140)
    Flags: SNMP-Traps 0x4000
    VLAN-Tag [0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530)
Out(swap-push 0x8100.512 0x8100.513)
    Encapsulation: VLAN-CCC
    Egress account overhead: 100
    Ingress account overhead: 90
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :                   0                    0 bps
     Output bytes  :                   0                    0 bps
     Input  packets:                   0                    0 pps
```

```
      Output packets:                    0                    0 pps
     Protocol ccc, MTU: 1522, Generation: 149, Route table: 0
       Flags: Is-Primary

   Logical interface ge-3/0/2.32767 (Index 71) (SNMP ifIndex 70)
   (Generation 139)
     Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
     Traffic statistics:
      Input  bytes  :                    0
      Output bytes  :                    0
      Input  packets:                    0
      Output packets:                    0
     Local statistics:
      Input  bytes  :                    0
      Output bytes  :                    0
      Input  packets:                    0
      Output packets:                    0
     Transit statistics:
      Input  bytes  :                    0                    0 bps
      Output bytes  :                    0                    0 bps
      Input  packets:                    0                    0 pps
      Output packets:                    0                    0 pps
```

show interfaces extensive (Gigabit Ethernet IQ2)

```
user@host> show interfaces ge-7/1/3 extensive
Physical interface: ge-7/1/3, Enabled, Physical link is Up
  Interface index: 170, SNMP ifIndex: 70, Generation: 171
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4004000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Schedulers     : 256
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:74, Hardware address: 00:00:5e:00:53:74
  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,
```

 Copyright © 2018, Juniper Networks, Inc.

```
  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Ingress queues: 4 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets

  0 best-effort           418390823            418390823                    0

  1 expedited-fo                  0                    0                    0

  2 assured-forw                  0                    0                    0

  3 network-cont               7133                 7133                    0

Egress queues: 4 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets

  0 best-effort                1031                 1031                    0

  1 expedited-fo                  0                    0                    0

  2 assured-forw                  0                    0                    0

  3 network-cont              77872                77872                    0

Active alarms  : None
Active defects : None
MAC statistics:                      Receive         Transmit
  Total octets                   38910844056          7174605
  Total packets                    418398473            78903
  Unicast packets               408021893366             1026
  Broadcast packets                       10               12
  Multicast packets                418398217            77865
  CRC/Align errors                         0                0
  FIFO errors                              0                0
  MAC control frames                       0                0
  MAC pause frames                         0                0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0  OTN Received Overhead Bytes:
  APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58
  Payload Type: 0x08
OTN Transmitted Overhead Bytes:
  APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
  Payload Type: 0x08
Filter statistics:
  Input packet count               418398473
  Input packet rejects                   479
  Input DA rejects                       479
  Input SA rejects                         0
  Output packet count                                   78903
  Output packet pad count                                   0
  Output packet error count                                 0
  CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Complete
  Link partner:
      Link mode: Full-duplex, Flow control: Symmetric/Asymmetric,
      Remote fault: OK
  Local resolution:
```

```
      Flow control: Symmetric, Remote fault: Link OK
  Packet Forwarding Engine configuration:
    Destination slot: 7
  CoS information:
    Direction : Output
    CoS transmit queue          Bandwidth          Buffer     Priority   Limit
                             %            bps     %           usec
    0 best-effort           95      950000000    95              0
low    none
    3 network-control        5       50000000     5              0
low    none
    Direction : Input
    CoS transmit queue          Bandwidth          Buffer     Priority   Limit
                             %            bps     %           usec
    0 best-effort           95      950000000    95              0
low    none
    3 network-control        5       50000000     5              0
low    none

  Logical interface ge-7/1/3.0 (Index 70) (SNMP ifIndex 85) (Generation 150)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :             812400
     Output bytes  :            1349206
     Input  packets:               9429
     Output packets:               9449
     IPv6 transit statistics:
      Input  bytes  :                0
      Output bytes  :                0
      Input  packets:                0
      Output packets:                0
    Local statistics:
     Input  bytes  :             812400
     Output bytes  :            1349206
     Input  packets:               9429
     Output packets:               9449
    Transit statistics:
     Input  bytes  :                 0               7440 bps
     Output bytes  :                 0               7888 bps
     Input  packets:                 0                 10 pps
     Output packets:                 0                 11 pps
     IPv6 transit statistics:
      Input  bytes  :                0
      Output bytes  :                0
      Input  packets:                0
      Output packets:                0
    Protocol inet, MTU: 1500, Generation: 169, Route table: 0
      Flags: Is-Primary, Mac-Validate-Strict
      Mac-Validate Failures: Packets: 0, Bytes: 0
      Addresses, Flags: Is-Preferred Is-Primary
      Input Filters: F1-ge-3/0/1.0-in, F3-ge-3/0/1.0-in
      Output Filters: F2-ge-3/0/1.0-out (53)
      Destination: 203.0.113/24, Local: 203.0.113.2, Broadcast: 203.0.113.255,
        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

 Copyright © 2018, Juniper Networks, Inc.

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 56 on page 1123.

show interfaces (Gigabit Ethernet Unnumbered Interface)

```
user@host> show interfaces ge-3/2/0
Physical interface: ge-3/2/0, Enabled, Physical link is Up
  Interface index: 148, SNMP ifIndex: 50
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Current address: 00:00:5e:00:53:f8, Hardware address: 00:00:5e:00:53: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: 203.0.113.22
```

show interfaces (ACI Interface Set Configured)

```
user@host> show interfaces ge-1/0/0.4001
  Logical interface ge-1/0/0.4001 (Index 340) (SNMP ifIndex 548)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.4001 ]  Encapsulation: PPP-over-

    Ethernet
    ACI VLAN:
      Dynamic Profile: aci-vlan-set-profile
    PPPoE:
      Dynamic Profile: aci-vlan-pppoe-profile,
      Service Name Table: None,
      Max Sessions: 32000, Max Sessions VSA Ignore: Off,
      Duplicate Protection: On, Short Cycle Protection: Off,
      Direct Connect: Off,
      AC Name: nbc
    Input packets : 9
    Output packets: 8
    Protocol multiservice, MTU: Unlimited
```

show interfaces (ALI Interface Set)

```
user@host> show interfaces ge-1/0/0.10
  Logical interface ge-1/0/0.10 (Index 346) (SNMP ifIndex 554) (Generation 155)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.10 ]  Encapsulation: ENET2
    Line Identity:
```

Copyright © 2018, Juniper Networks, Inc.

```
      Dynamic Profile: ali-set-profile
      Circuit-id Remote-id Accept-no-ids
    PPPoE:
      Dynamic Profile: ali-vlan-pppoe-profile,
      Service Name Table: None,
      Max Sessions: 32000, Max Sessions VSA Ignore: Off,
      Duplicate Protection: On, Short Cycle Protection: Off,
      Direct Connect: Off,
      AC Name: nbc
    Input packets : 9
    Output packets: 8
    Protocol multiservice, MTU: Unlimited
```

Sample Output Gigabit Ethernet

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, IQ2)

```
user@host> show interfaces xe-5/0/0 extensive
Physical interface: xe-5/0/0, Enabled, Physical link is Up
  Interface index: 177, SNMP ifIndex: 99, Generation: 178
  Link-level type: Ethernet, MTU: 1518, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Enabled,
  Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Schedulers     : 1024
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:f6, Hardware address: 00:00:5e:00:53: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
```

 Copyright © 2018, Juniper Networks, Inc.

```
  2 assured-forw                  0                    0                    0

  3 network-cont                  0                    0                    0

Egress queues: 4 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets

  0 best-effort                   0                    0                    0

  1 expedited-fo                  0                    0                    0

  2 assured-forw                  0                    0                    0

  3 network-cont                  0                    0                    0

Active alarms  : None
Active defects : None
PCS statistics                     Seconds
  Bit errors                           0
  Errored blocks                       0
MAC statistics:                      Receive         Transmit
  Total octets                    6970332384                0
  Total packets                     81050506                0
  Unicast packets                   81050000                0
  Broadcast packets                      506                0
  Multicast packets                        0                0
  CRC/Align errors                         0                0
  FIFO errors                              0                0
  MAC control frames                       0                0
  MAC pause frames                         0                0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0
Filter statistics:
  Input packet count                81050506
  Input packet rejects                   506
  Input DA rejects                         0
  Input SA rejects                         0
  Output packet count                                       0
  Output packet pad count                                   0
  Output packet error count                                 0
  CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 5
CoS information:
  Direction : Output
  CoS transmit queue          Bandwidth               Buffer Priority   Limit
                       %            bps     %         usec
  0 best-effort       95      950000000    95         0           low     none
  3 network-control    5       50000000     5         0           low     none

  Direction : Input
  CoS transmit queue          Bandwidth               Buffer Priority   Limit
                       %            bps     %         usec
  0 best-effort       95      950000000    95         0           low     none
  3 network-control    5       50000000     5         0           low     none

Logical interface xe-5/0/0.0 (Index 71) (SNMP ifIndex 95) (Generation 195)
```

```
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ]  Encapsulation: ENET2
    Egress account overhead: 100
    Ingress account overhead: 90
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                  46
     Input  packets:                   0
     Output packets:                   1
     IPv6 transit statistics:
      Input  bytes  :                  0
      Output bytes  :                  0
      Input  packets:                  0
      Output packets:                  0
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                  46
     Input  packets:                   0
     Output packets:                   1
    Transit statistics:
     Input  bytes  :                   0                    0 bps
     Output bytes  :                   0                    0 bps
     Input  packets:                   0                    0 pps
     Output packets:                   0                    0 pps
     IPv6 transit statistics:
      Input  bytes  :                  0
      Output bytes  :                  0
      Input  packets:                  0
      Output packets:                  0
    Protocol inet, MTU: 1500, Generation: 253, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.0.2/24, Local: 192.0.2.1, Broadcast: 192.0.2.255,
Generation: 265
    Protocol multiservice, MTU: Unlimited, Generation: 254, Route table: 0
      Flags: None
      Policer: Input: __default_arp_policer__
```

show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode)

```
user@host> show interfaces xe-1/0/0 extensive
Physical interface: xe-1/0/0, Enabled, Physical link is Up
  Interface index: 141, SNMP ifIndex: 34, Generation: 47
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Loopback: Disabled
  WAN-PHY mode
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps 16384
  Link flags     : None
  CoS queues     : 4 supported
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:9d, Hardware address: 00:00:5e:00:53:9d
  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,
```

 Copyright © 2018, Juniper Networks, Inc.

```
      Resource errors: 0
    Output errors:
      Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0,
      Aged packets: 0, FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0,
      Resource errors: 0
    Queue counters:       Queued packets  Transmitted packets     Dropped packets
      0 best-effort                    0                    0                   0
      1 expedited-fo                   0                    0                   0
      2 assured-forw                   0                    0                   0
      3 network-cont                   0                    0                   0
    Active alarms : LOL, LOS, LBL
    Active defects: LOL, LOS, LBL, SEF, AIS-L, AIS-P
    PCS statistics             Seconds       Count
      Bit errors                     0           0
      Errored blocks                 0           0
    MAC statistics:                      Receive         Transmit
      Total octets                             0                0
      Total packets                            0                0
      Unicast packets                          0                0
      Broadcast packets                        0                0
      Multicast packets                        0                0
      CRC/Align errors                         0                0
      FIFO errors                              0                0
      MAC control frames                       0                0
      MAC pause frames                         0                0
      Oversized frames                         0
      Jabber frames                            0
      Fragment frames                          0
      VLAN tagged frames                       0
      Code violations                          0
    Filter statistics:
      Input packet count                       0
      Input packet rejects                     0
      Input DA rejects                         0
      Input SA rejects                         0
      Output packet count                                       0
      Output packet pad count                                   0
      Output packet error count                                 0
      CAM destination filters: 0, CAM source filters: 0
    PMA PHY:              Seconds        Count  State
      PLL lock                  0            0  OK
      PHY light             63159            1  Light Missing
   WIS section:
     BIP-B1                     0           0
     SEF                   434430      434438  Defect Active
     LOS                   434430           1  Defect Active
     LOF                   434430           1  Defect Active
     ES-S                  434430
     SES-S                 434430
     SEFS-S                434430
   WIS line:
     BIP-B2                     0           0
     REI-L                      0           0
     RDI-L                      0           0  OK
     AIS-L                 434430           1  Defect Active
     BERR-SF                    0           0  OK
     BERR-SD                    0           0  OK
     ES-L                  434430
     SES-L                 434430
     UAS-L                 434420
     ES-LFE                     0
```

```
  SES-LFE                 0
  UAS-LFE                 0
WIS path:
  BIP-B3                  0            0
  REI-P                   0            0
  LOP-P                   0            0  OK
  AIS-P              434430            1  Defect Active
  RDI-P                   0            0  OK
  UNEQ-P                  0            0  OK
  PLM-P                   0            0  OK
  ES-P               434430
  SES-P              434430
  UAS-P              434420
  ES-PFE                  0
  SES-PFE                 0
  UAS-PFE                 0
Received path trace:
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
Transmitted path trace: orissa so-1/0/0
  6f 72 69 73 73 61 20 73 6f 2d 31 2f 30 2f 30 00   orissa so-1/0/0.
  Packet Forwarding Engine configuration:
    Destination slot: 1
  CoS information:
    CoS transmit queue          Bandwidth          Buffer     Priority   Limit
                          %           bps     %    bytes
    0 best-effort        95     950000000    95        0          low     none
    3 network-control     5      50000000     5        0          low     none
```

show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC)

```
user@host> show interfaces ge-7/0/0 extensive
Physical interface: ge-7/0/0, Enabled, Physical link is Down
  Interface index: 143, SNMP ifIndex: 508, Generation: 208
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, BPDU Error: None,
  MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Enabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
  Link flags     : None
  Wavelength     : 1550.12 nm, Frequency: 193.40 THz
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:72, Hardware address: 00:00:5e:00:53:72
  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,
```

 Copyright © 2018, Juniper Networks, Inc.

```
  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
```

Copyright © 2018, Juniper Networks, Inc.

```
    LOM                        67164          71  Defect Active
    Wavelength Lock                0           0  OK
  OTN OTU:
    AIS                            0           0  OK
    BDI                        65919        4814  Defect Active
    IAE                        67158           1  Defect Active
    TTIM                           7           1  OK
    SF                         67164           2  Defect Active
    SD                         67164           3  Defect Active
    TCA-ES                         0           0  OK
    TCA-SES                        0           0  OK
    TCA-UAS                       80          40  OK
    TCA-BBE                        0           0  OK
    BIP                            0           0  OK
    BBE                            0           0  OK
    ES                             0           0  OK
    SES                            0           0  OK
    UAS                          587           0  OK
    Received DAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Received SAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Transmitted DAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Transmitted SAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
  OTN Received Overhead Bytes:
    APS/PCC0: 0x02, APS/PCC1: 0x42, APS/PCC2: 0xa2, APS/PCC3: 0x48
    Payload Type: 0x03
  OTN Transmitted Overhead Bytes:
    APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
    Payload Type: 0x03
  Filter statistics:
    Input packet count                        0
    Input packet rejects                      0
    Input DA rejects                          0
    Input SA rejects                          0
    Output packet count                                        0
    Output packet pad count                                    0
    Output packet error count                                  0
    CAM destination filters: 0, CAM source filters: 0
  Packet Forwarding Engine configuration:
    Destination slot: 7
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority
Limit
                              %            bps     %           usec
    0 best-effort            95     9500000000    95              0       low
none
    3 network-control         5      500000000     5              0       low
none
    ...
```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode)

```
user@host> show interfaces xe-7/0/0 extensive
Physical interface: xe-7/0/0, Enabled, Physical link is Up
  Interface index: 173, SNMP ifIndex: 212, Generation: 174
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
Unidirectional: Enabled,
```

 Copyright © 2018, Juniper Networks, Inc.

```
  Loopback: None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
...
```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only)

```
user@host> show interfaces xe-7/0/0–tx extensive
Physical interface: xe-7/0/0-tx, Enabled, Physical link is Up
  Interface index: 176, SNMP ifIndex: 137, Generation: 177
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
Unidirectional: Tx-Only
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:83, Hardware address: 00:00:5e:00:53:83
  Last flapped   : 2007-06-01 09:08:19 PDT (3d 02:31 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :      322891152287160           9627472888 bps
   Input  packets:                    0                    0 pps
   Output packets:         328809727380              1225492 pps

...

  Filter statistics:
    Output packet count           328810554250
    Output packet pad count                  0
    Output packet error count                0
...

  Logical interface xe-7/0/0-tx.0 (Index 73) (SNMP ifIndex 138) (Generation 139)

    Flags: SNMP-Traps Encapsulation: ENET2
    Egress account overhead: 100
    Ingress account overhead: 90
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :      322891152287160
     Input  packets:                    0
     Output packets:         328809727380
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :      322891152287160           9627472888 bps
     Input  packets:                    0                    0 pps
     Output packets:         328809727380              1225492 pps
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
```

```
      Input  packets:                   0
      Output packets:                   0
    Protocol inet, MTU: 1500, Generation: 147, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.11.12/24, Local: 10.11.12.13, Broadcast: 10.11.12.255,
Generation: 141
    Protocol multiservice, MTU: Unlimited, Generation: 148, Route table: 0
      Flags: None
      Policer: Input: __default_arp_policer__
```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only)

```
user@host> show interfaces xe-7/0/0–rx extensive
Physical interface: xe-7/0/0-rx, Enabled, Physical link is Up
  Interface index: 174, SNMP ifIndex: 118, Generation: 175
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
Unidirectional: Rx-Only
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:83, Hardware address: 00:00:5e:00:53:83
  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
```

 Copyright © 2018, Juniper Networks, Inc.

```
     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: 192.0.2/24, Local: 192.0.2.1, Broadcast: 192.0.2.255,
Generation: 139
    Protocol multiservice, MTU: Unlimited, Generation: 146, Route table: 0
      Flags: None
      Policer: Input: __default_arp_policer__
```

Sample Output

Sample Output SRX Gigabit Ethernet

```
user@host> show interfaces ge-0/0/1
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, 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 Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:42 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK
  Interface transmit statistics: Disabled

  Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514)
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Security: Zone: public
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255
```

Sample Output SRX Gigabit Ethernet

```
user@host> show interfaces ge-0/0/1
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, 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 Down
```

```
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:42 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK
  Interface transmit statistics: Disabled

  Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514)
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Security: Zone: public
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255
```

show interfaces detail (Gigabit Ethernet)

user@host> show interfaces ge-0/0/1 detail

```
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510, Generation: 138
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, 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 Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w2d 00:00 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
  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  : LINK
  Active defects : LINK
```


Copyright © 2018, Juniper Networks, Inc.

```
 Interface transmit statistics: Disabled

 Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)
   Flags: Device-Down SNMP-Traps 0x0 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
   Security: Zone: public
   Flow Statistics :
   Flow Input statistics :
     Self packets :                    0
     ICMP packets :                    0
     VPN packets :                     0
     Multicast packets :               0
     Bytes permitted by policy :       0
     Connections established :         0
   Flow Output statistics:
     Multicast packets :               0
     Bytes permitted by policy :       0
   Flow error statistics (Packets dropped due to):
     Address spoofing:                 0
     Authentication failed:            0
     Incoming NAT errors:              0
     Invalid zone received packet:     0
     Multiple user authentications:    0
     Multiple incoming NAT:            0
     No parent for a gate:             0
     No one interested in self packets: 0
     No minor session:                 0
     No more sessions:                 0
     No NAT gate:                      0
     No route present:                 0
     No SA for incoming SPI:           0
     No tunnel found:                  0
     No session for a gate:            0
     No zone or NULL zone binding      0
     Policy denied:                    0
     Security association not active:  0
     TCP sequence number out of window: 0
     Syn-attack protection:            0
     User authentication errors:       0
   Protocol inet, MTU: 1500, Generation: 150, Route table: 0
     Flags: Sendbcast-pkt-to-re
     Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
       Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255, Generation:
150
```

show interfaces statistics st0.0 detail

```
user@host> show interfaces statistics st0.0 detail
  Logical interface st0.0 (Index 71) (SNMP ifIndex 609) (Generation 136)
    Flags: Up Point-To-Point SNMP-Traps Encapsulation: Secure-Tunnel
    Traffic statistics:
     Input  bytes  :        528152756774
     Output bytes  :        575950643520
     Input  packets:         11481581669
     Output packets:         12520666095
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :                   0            121859888 bps
     Output bytes  :                   0            128104112 bps
     Input  packets:                   0               331141 pps
     Output packets:                   0               348108 pps
    Security: Zone: untrust
    Allowed host-inbound traffic : any-service bfd bgp dvmrp igmp ldp msdp nhrp
ospf ospf3 pgm pim rip ripng router-discovery rsvp
    sap vrrp
    Flow Statistics :
    Flow Input statistics :
      Self packets :                     0
      ICMP packets :                     0
      VPN packets :                      0
      Multicast packets :                0
      Bytes permitted by policy :        525984295844
      Connections established :          7
    Flow Output statistics:
      Multicast packets :                0
      Bytes permitted by policy :        576003290222
    Flow error statistics (Packets dropped due to):
      Address spoofing:                  0
      Authentication failed:             0
      Incoming NAT errors:               0
      Invalid zone received packet:      0
      Multiple user authentications:     0
      Multiple incoming NAT:             0
      No parent for a gate:              0
      No one interested in self packets: 0
      No minor session:                  0
      No more sessions:                  0
      No NAT gate:                       0
      No route present:                  2000280
      No SA for incoming SPI:            0
      No tunnel found:                   0
      No session for a gate:             0
      No zone or NULL zone binding       0
      Policy denied:                     0
      Security association not active:   0
      TCP sequence number out of window: 0
      Syn-attack protection:             0
      User authentication errors:        0
    Protocol inet, MTU: 9192
    Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0,
 NH drop cnt: 0
    Generation: 155, Route table: 0
```


Copyright © 2018, Juniper Networks, Inc.

```
        Flags: Sendbcast-pkt-to-re
```

show interfaces extensive (Gigabit Ethernet)

```
user@host> show interfaces ge-0/0/1.0 extensive
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510, Generation: 138
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, 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 Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:57 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,
    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  : 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
```

Copyright © 2018, Juniper Networks, Inc.

```
    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: 2, CAM source filters: 0
  Autonegotiation information:
    Negotiation status: Incomplete
  Packet Forwarding Engine configuration:
    Destination slot: 0
  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
  Interface transmit statistics: Disabled

  Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)
    Flags: Device-Down SNMP-Traps 0x0 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
    Security: Zone: public
    Flow Statistics :
    Flow Input statistics :
      Self packets :                  0
      ICMP packets :                  0
      VPN packets :                   0
      Multicast packets :             0
      Bytes permitted by policy :     0
      Connections established :       0
    Flow Output statistics:
      Multicast packets :             0
      Bytes permitted by policy :     0
    Flow error statistics (Packets dropped due to):
      Address spoofing:               0
```

 Copyright © 2018, Juniper Networks, Inc.

```
    Authentication failed:            0
    Incoming NAT errors:              0
    Invalid zone received packet:     0
    Multiple user authentications:    0
    Multiple incoming NAT:            0
    No parent for a gate:             0
    No one interested in self packets: 0
    No minor session:                 0
    No more sessions:                 0
    No NAT gate:                      0
    No route present:                 0
    No SA for incoming SPI:           0
    No tunnel found:                  0
    No session for a gate:            0
    No zone or NULL zone binding      0
    Policy denied:                    0
    Security association not active:  0
    TCP sequence number out of window: 0
    Syn-attack protection:            0
    User authentication errors:       0
  Protocol inet, MTU: 1500, Generation: 150, Route table: 0
    Flags: Sendbcast-pkt-to-re
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255,
      Generation: 150
```

show interfaces terse

```
user@host> show interfaces terse
    Interface               Admin Link Proto    Local                 Remote
    ge-0/0/0                up    up
    ge-0/0/0.0              up    up   inet     10.209.4.61/18
    gr-0/0/0                up    up
    ip-0/0/0                up    up
    st0                     up    up
    st0.1                   up    ready inet
    ls-0/0/0                up    up
    lt-0/0/0                up    up
    mt-0/0/0                up    up
    pd-0/0/0                up    up
    pe-0/0/0                up    up
    e3-1/0/0                up    up
    t3-2/0/0                up    up
    e1-3/0/0                up    up
    se-4/0/0                up    down
    t1-5/0/0                up    up
    br-6/0/0                up    up
    dc-6/0/0                up    up
    dc-6/0/0.32767          up    up
    bc-6/0/0:1              down  up
    bc-6/0/0:1.0            up    down
    dl0                     up    up
    dl0.0                   up    up   inet
    dsc                     up    up
    gre                     up    up
    ipip                    up    up
    lo0                     up    up
    lo0.16385               up    up   inet     10.0.0.1            --> 0/0
                                                10.0.0.16           --> 0/0
```

```
lsi                     up    up
mtun                    up    up
pimd                    up    up
pime                    up    up
pp0                     up    up
```

show interfaces controller (Channelized E1 IQ with Logical E1)

```
user@host> show interfaces controller ce1-1/2/6

Controller                                                          Admin Link
ce1-1/2/6                                                           up    up
e1-1/2/6                                                            up    up
```

show interfaces controller (Channelized E1 IQ with Logical DS0)

```
user@host> show interfaces controller ce1-1/2/3

Controller                                                          Admin Link
ce1-1/2/3                                                           up    up

ds-1/2/3:1                                                          up    up

ds-1/2/3:2                                                          up    up
```

show interfaces descriptions

```
user@host> show interfaces descriptions
 Interface       Admin Link Description
 so-1/0/0        up    up   M20-3#1
 so-2/0/0        up    up   GSR-12#1
 ge-3/0/0        up    up   SMB-OSPF_Area300
 so-3/3/0        up    up   GSR-13#1
 so-3/3/1        up    up   GSR-13#2
 ge-4/0/0        up    up   T320-7#1
 ge-5/0/0        up    up   T320-7#2
 so-7/1/0        up    up   M160-6#1
 ge-8/0/0        up    up   T320-7#3
 ge-9/0/0        up    up   T320-7#4
 so-10/0/0       up    up   M160-6#2
 so-13/0/0       up    up   M20-3#2
 so-14/0/0       up    up   GSR-12#2
 ge-15/0/0       up    up   SMB-OSPF_Area100
 ge-15/0/1       up    up   GSR-13#3
```

show interfaces destination-class all

```
user@host> show interfaces destination-class all
Logical interface so-4/0/0.0
                                          Packets                Bytes
      Destination class            (packet-per-second)    (bits-per-second)
                         gold                    0                    0
                                  (              0) (                 0)
                       silver                    0                    0
                                  (              0) (                 0)
  Logical interface so-0/1/3.0
                                          Packets                Bytes
      Destination class            (packet-per-second)    (bits-per-second)
                         gold                    0                    0
                                  (              0) (                 0)
```

 Copyright © 2018, Juniper Networks, Inc.

```
                        silver                        0                    0
                                   (                  0) (                 0)
```

show interfaces diagnostics optics

```
user@host> show interfaces diagnostics optics ge-2/0/0
Physical interface: ge-2/0/0
    Laser bias current                        :  7.408 mA
    Laser output power                        :  0.3500 mW / -4.56 dBm
    Module temperature                        :  23 degrees C / 73 degrees F
    Module voltage                            :  3.3450 V
    Receiver signal average optical power     :  0.0002 mW / -36.99 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  On
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  On
    Laser bias current high alarm threshold   :  17.000 mA
    Laser bias current low alarm threshold    :  1.000 mA
    Laser bias current high warning threshold :  14.000 mA
    Laser bias current low warning threshold  :  2.000 mA
    Laser output power high alarm threshold   :  0.6310 mW / -2.00 dBm
    Laser output power low alarm threshold    :  0.0670 mW / -11.74 dBm
    Laser output power high warning threshold :  0.6310 mW / -2.00 dBm
    Laser output power low warning threshold  :  0.0790 mW / -11.02 dBm
    Module temperature high alarm threshold   :  95 degrees C / 203 degrees F
    Module temperature low alarm threshold    :  -25 degrees C / -13 degrees F
    Module temperature high warning threshold :  90 degrees C / 194 degrees F
    Module temperature low warning threshold  :  -20 degrees C / -4 degrees F
    Module voltage high alarm threshold       :  3.900 V
    Module voltage low alarm threshold        :  2.700 V
    Module voltage high warning threshold     :  3.700 V
    Module voltage low warning threshold      :  2.900 V
    Laser rx power high alarm threshold       :  1.2590 mW / 1.00 dBm
    Laser rx power low alarm threshold        :  0.0100 mW / -20.00 dBm
    Laser rx power high warning threshold     :  0.7940 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.0158 mW / -18.01 dBm
```

show interfaces far-end-interval coc12-5/2/0

```
user@host> show interfaces far-end-interval coc12-5/2/0
Physical interface: coc12-5/2/0, SNMP ifIndex: 121
  05:30-current:
    ES-L: 1, SES-L: 1, UAS-L: 0
```

```
  05:15-05:30:
    ES-L: 0, SES-L: 0, UAS-L: 0
  05:00-05:15:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:45-05:00:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:30-04:45:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:15-04:30:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:00-04:15:
...
```

show interfaces far-end-interval coc1-5/2/1:1

```
user@host> run show interfaces far-end-interval coc1-5/2/1:1
 Physical interface: coc1-5/2/1:1, SNMP ifIndex: 342
  05:30-current:
    ES-L: 1, SES-L: 1, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  05:15-05:30:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  05:00-05:15:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:45-05:00:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:30-04:45:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:15-04:30:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:00-04:15:
```

show interfaces filters

```
user@host> show interfaces filters
Interface       Admin Link Proto Input Filter         Output Filter
ge-0/0/0        up    up
ge-0/0/0.0      up    up   inet
                           iso
ge-5/0/0        up    up
ge-5/0/0.0      up    up   any                        f-any
                           inet                       f-inet
                           multiservice
gr-0/3/0        up    up
ip-0/3/0        up    up
mt-0/3/0        up    up
pd-0/3/0        up    up
pe-0/3/0        up    up
vt-0/3/0        up    up
at-1/0/0        up    up
at-1/0/0.0      up    up   inet
                           iso
at-1/1/0        up    down
at-1/1/0.0      up    down inet
                           iso
....
```

show interfaces flow-statistics (Gigabit Ethernet)

```
user@host> show interfaces flow-statistics ge-0/0/1.0
```

 Copyright © 2018, Juniper Networks, Inc.

```
  Logical interface ge-0/0/1.0 (Index 70) (SNMP ifIndex 49)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 5161
    Output packets: 83
    Security: Zone: zone2
    Allowed host-inbound traffic : bootp bfd bgp  dns dvmrp  ldp msdp nhrp ospf
pgm
    pim rip router-discovery rsvp sap vrrp dhcp finger ftp tftp ident-reset http
 https ike
   netconf ping rlogin rpm rsh snmp snmp-trap ssh telnet traceroute xnm-clear-text
 xnm-ssl
    lsping
    Flow Statistics :
    Flow Input statistics :
      Self packets :                     0
      ICMP packets :                     0
      VPN packets :                      2564
      Bytes permitted by policy :        3478
      Connections established :          1
    Flow Output statistics:
      Multicast packets :                0
      Bytes permitted by policy :        16994
    Flow error statistics (Packets dropped due to):
      Address spoofing:                  0
      Authentication failed:             0
      Incoming NAT errors:               0
      Invalid zone received packet:      0
      Multiple user authentications:     0
      Multiple incoming NAT:             0
      No parent for a gate:              0
      No one interested in self packets: 0
      No minor session:                  0
      No more sessions:                  0
      No NAT gate:                       0
      No route present:                  0
      No SA for incoming SPI:            0
      No tunnel found:                   0
      No session for a gate:             0
      No zone or NULL zone binding       0
      Policy denied:                     0
      Security association not active:   0
      TCP sequence number out of window: 0
      Syn-attack protection:             0
      User authentication errors:        0
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination:  203.0.113.1/24, Local: 203.0.113.2, Broadcast: 2.2.2.255
```

show interfaces interval (Channelized OC12)

```
user@host> show interfaces interval t3-0/3/0:0
Physical interface: t3-0/3/0:0, SNMP ifIndex: 23
  17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:28-17:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:13-17:28:
```

Copyright © 2018, Juniper Networks, Inc.

```
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:58-17:13:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
...
Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
    CES: 230, CSES: 230, SEFS: 230, UAS: 238
```

show interfaces interval (E3)

```
user@host> show interfaces interval e3-0/3/0
Physical interface: e3-0/3/0, SNMP ifIndex: 23
  17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:28-17:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:13-17:28:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:58-17:13:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
....
Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
    CES: 230, CSES: 230, SEFS: 230, UAS: 238
```

show interfaces interval (SONET/SDH) (SRX devices)

```
user@host> show interfaces interval so-0/1/0
Physical interface: so-0/1/0, SNMP ifIndex: 19
  20:02-current:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0,
    SES-P: 0, UAS-P: 0
  19:47-20:02:
    ES-S: 267, SES-S: 267, SEFS-S: 267, ES-L: 267, SES-L: 267, UAS-L: 267,
    ES-P: 267, SES-P: 267, UAS-P: 267
  19:32-19:47:
    ES-S: 56, SES-S: 56, SEFS-S: 56, ES-L: 56, SES-L: 56, UAS-L: 46, ES-P: 56,
    SES-P: 56, UAS-P: 46
  19:17-19:32:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0,
    SES-P: 0, UAS-P: 0
  19:02-19:17:
.....
```

show interfaces load-balancing (SRX devices)

```
user@host> show interfaces load-balancing
Interface  State            Last change  Member count
ams0       Up               1d 00:50     2
ams1       Up               00:00:59     2
```

 Copyright © 2018, Juniper Networks, Inc.

show interfaces load-balancing detail (SRX devices)

```
user@host>show interfaces load-balancing detail
Load-balancing interfaces detail
Interface         : ams0
  State           : Up
  Last change     : 1d 00:51
  Member count    : 2
  Members         :
      Interface     Weight    State
      mams-2/0/0     10         Active
      mams-2/1/0     10         Active
```

show interfaces mac-database (All MAC Addresses on a Port SRX devices)

```
user@host> show interfaces mac-database xe-0/3/3
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags    : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags      : None

  Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address         Input frames     Input bytes  Output frames    Output bytes
  00:00:00:00:00:00              1              56              0               0
  00:00:c0:01:01:02        7023810       323095260              0               0
  00:00:c0:01:01:03        7023810       323095260              0               0
  00:00:c0:01:01:04        7023810       323095260              0               0
  00:00:c0:01:01:05        7023810       323095260              0               0
  00:00:c0:01:01:06        7023810       323095260              0               0
  00:00:c0:01:01:07        7023810       323095260              0               0
  00:00:c0:01:01:08        7023809       323095214              0               0
  00:00:c0:01:01:09        7023809       323095214              0               0
  00:00:c0:01:01:0a        7023809       323095214              0               0
  00:00:c0:01:01:0b        7023809       323095214              0               0
  00:00:c8:01:01:02       30424784      1399540064       37448598      1722635508
  00:00:c8:01:01:03       30424784      1399540064       37448598      1722635508
  00:00:c8:01:01:04       30424716      1399536936       37448523      1722632058
  00:00:c8:01:01:05       30424789      1399540294       37448598      1722635508
  00:00:c8:01:01:06       30424788      1399540248       37448597      1722635462
  00:00:c8:01:01:07       30424783      1399540018       37448597      1722635462
  00:00:c8:01:01:08       30424783      1399540018       37448596      1722635416
  00:00:c8:01:01:09        8836796       406492616        8836795       406492570
  00:00:c8:01:01:0a       30424712      1399536752       37448521      1722631966
  00:00:c8:01:01:0b       30424715      1399536890       37448523      1722632058
  Number of MAC addresses : 21
```

show interfaces mac-database (All MAC Addresses on a Service SRX devices)

```
user@host> show interfaces mac-database xe-0/3/3
Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address         Input frames     Input bytes  Output frames    Output bytes
  00:00:00:00:00:00              1              56              0               0
```

```
00:00:c0:01:01:02       7023810     323095260           0           0
00:00:c0:01:01:03       7023810     323095260           0           0
00:00:c0:01:01:04       7023810     323095260           0           0
00:00:c0:01:01:05       7023810     323095260           0           0
00:00:c0:01:01:06       7023810     323095260           0           0
00:00:c0:01:01:07       7023810     323095260           0           0
00:00:c0:01:01:08       7023809     323095214           0           0
00:00:c0:01:01:09       7023809     323095214           0           0
00:00:c0:01:01:0a       7023809     323095214           0           0
00:00:c0:01:01:0b       7023809     323095214           0           0
00:00:c8:01:01:02      31016568    1426762128    38040381  1749857526
00:00:c8:01:01:03      31016568    1426762128    38040382  1749857572
00:00:c8:01:01:04      31016499    1426758954    38040306  1749854076
00:00:c8:01:01:05      31016573    1426762358    38040381  1749857526
00:00:c8:01:01:06      31016573    1426762358    38040381  1749857526
00:00:c8:01:01:07      31016567    1426762082    38040380  1749857480
00:00:c8:01:01:08      31016567    1426762082    38040379  1749857434
00:00:c8:01:01:09       9428580     433714680     9428580   433714680
00:00:c8:01:01:0a      31016496    1426758816    38040304  1749853984
00:00:c8:01:01:0b      31016498    1426758908    38040307  1749854122
```

show interfaces mac-database mac-address

```
user@host> show interfaces mac-database xe-0/3/3 mac-address (SRX devices)
00:00:c8:01:01:09
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address: 00:00:c8:01:01:09, Type: Configured,
    Input bytes     : 202324652
    Output bytes    : 202324560
    Input frames    : 4398362
    Output frames   : 4398360
  Policer statistics:
  Policer type      Discarded frames    Discarded bytes
  Output aggregate           3992386          183649756
```

show interfaces mc-ae (SRX devices)

```
user@host> show interfaces mc-ae ae0 unit 512
  Member Links   : ae0
  Local Status   : active
  Peer Status    : active
  Logical Interface        : ae0.512
  Core Facing Interface : Label Ethernet Interface
  ICL-PL         : Label Ethernet Interface
```

show interfaces media (SONET/SDH)

The following example displays the output fields unique to the **show interfaces media** command for a SONET interface (with no level of output specified):

 Copyright © 2018, Juniper Networks, Inc.

```
user@host> show interfaces media so-4/1/2
Physical interface: so-4/1/2, Enabled, Physical link is Up
  Interface index: 168, SNMP ifIndex: 495
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 16384
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 1783 (00:00:00 ago), Output: 1786 (00:00:08 ago)
  LCP state: Opened
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
  mpls: Not-configured
  CHAP state: Not-configured
  CoS queues     : 8 supported
  Last flapped   : 2005-06-15 12:14:59 PDT (04:31:29 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : None
  SONET defects  : None
  SONET errors:
     BIP-B1: 121, BIP-B2: 916, REI-L: 0, BIP-B3: 137, REI-P: 16747, BIP-BIP2: 0
  Received path trace: routerb so-1/1/2
  Transmitted path trace: routera so-4/1/2
```

show interfaces policers (SRX devices)

```
user@host> show interfaces policers
Interface       Admin Link Proto Input Policer         Output Policer
ge-0/0/0        up    up
ge-0/0/0.0      up    up   inet
                           iso
gr-0/3/0        up    up
ip-0/3/0        up    up
mt-0/3/0        up    up
pd-0/3/0        up    up
pe-0/3/0        up    up
...
so-2/0/0        up    up
so-2/0/0.0      up    up   inet  so-2/0/0.0-in-policer so-2/0/0.0-out-policer
                           iso
so-2/1/0        up    down
...
```

show interfaces policers interface-name (SRX devices)

```
user@host> show interfaces policers so-2/1/0
Interface       Admin Link Proto Input Policer         Output Policer
so-2/1/0        up    down
so-2/1/0.0      up    down inet  so-2/1/0.0-in-policer so-2/1/0.0-out-policer
                           iso
                           inet6
```

show interfaces queue (SRX devices)

The following truncated example shows the CoS queue sizes for queues 0, 1, and 3. Queue 1 has a queue buffer size (guaranteed allocated memory) of 9192 bytes.

```
user@host> show interfaces queue
Physical interface: ge-0/0/0, Enabled, Physical link is Up
  Interface index: 134, SNMP ifIndex: 509
Forwarding classes: 8 supported, 8 in use
Egress queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: class0
  Queued:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
  Transmitted:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
    Tail-dropped packets :                     0                     0 pps
    RL-dropped packets   :                     0                     0 pps
    RL-dropped bytes     :                     0                     0 bps
    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 Buffer Usage:
    Reserved buffer      :             118750000 bytes
  Queue-depth bytes      :
    Current              :                     0
..
..
Queue: 1, Forwarding classes: class1
..
..
 Queue Buffer Usage:
    Reserved buffer      :                  9192 bytes
  Queue-depth bytes      :
    Current              :                     0
..
..
Queue: 3, Forwarding classes: class3
  Queued:
..
..
Queue Buffer Usage:
    Reserved buffer      :               6250000 bytes
  Queue-depth bytes      :
    Current              :                     0
..
..
```

show interfaces redundancy (SRX devices)

```
user@host> show interfaces redundancy
Interface  State         Last change  Primary    Secondary  Current status
rsp0       Not present                sp-1/0/0   sp-0/2/0   both down
rsp1       On secondary  1d 23:56     sp-1/2/0   sp-0/3/0   primary down
rsp2       On primary    10:10:27     sp-1/3/0   sp-0/2/0   secondary down
rlsq0      On primary    00:06:24     lsq-0/3/0  lsq-1/0/0  both up
```

 Copyright © 2018, Juniper Networks, Inc.

show interfaces redundancy (Aggregated Ethernet SRX devices)

```
user@host> show interfaces redundancy
Interface   State           Last change  Primary      Secondary    Current status
rlsq0       On secondary     00:56:12     lsq-4/0/0    lsq-3/0/0    both up

ae0
ae1
ae2
ae3
ae4
```

show interfaces redundancy detail (SRX devices)

```
user@host> show interfaces redundancy detail
Interface        : rlsq0
  State          : On primary
  Last change    : 00:45:47
  Primary        : lsq-0/2/0
  Secondary      : lsq-1/2/0
  Current status : both up
  Mode           : hot-standby

Interface        : rlsq0:0
  State          : On primary
  Last change    : 00:45:46
  Primary        : lsq-0/2/0:0
  Secondary      : lsq-1/2/0:0
  Current status : both up
  Mode           : warm-standby
```

show interfaces routing brief (SRX devices)

```
user@host> show interfaces routing brief
Interface        State Addresses
so-5/0/3.0       Down  ISO   enabled
so-5/0/2.0       Up    MPLS  enabled
                       ISO   enabled
                       INET  192.168.2.120
                       INET  enabled
so-5/0/1.0       Up    MPLS  enabled
                       ISO   enabled
                       INET  192.168.2.130
                       INET  enabled
at-1/0/0.3       Up    CCC   enabled
at-1/0/0.2       Up    CCC   enabled
at-1/0/0.0       Up    ISO   enabled
                       INET  192.168.90.10
                       INET  enabled
lo0.0            Up    ISO   47.0005.80ff.f800.0000.0108.0001.1921.6800.5061.00
                       ISO   enabled
                       INET  127.0.0.1
fxp1.0           Up
fxp0.0           Up    INET  192.168.6.90
```

show interfaces routing detail (SRX devices)

```
user@host> show interfaces routing detail
so-5/0/3.0
  Index: 15, Refcount: 2, State: Up <Broadcast PointToPoint Multicast> Change:<>
```

```
  Metric: 0, Up/down transitions: 0, Full-duplex
  Link layer: HDLC serial line Encapsulation: PPP Bandwidth: 155Mbps
  ISO address (null)
    State: <Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
so-5/0/2.0
  Index: 14, Refcount: 7, State: <Up Broadcast PointToPoint Multicast> Change:<>

  Metric: 0, Up/down transitions: 0, Full-duplex
  Link layer: HDLC serial line Encapsulation: PPP Bandwidth: 155Mbps
  MPLS address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4458 bytes
  ISO address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
  INET address 192.168.2.120
    State: <Up Broadcast PointToPoint Multicast Localup> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
    Local address: 192.168.2.120
    Destination: 192.168.2.110/32
  INET address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
...
```

show interfaces routing-instance all (SRX devices)

```
user@host> show interfaces terse routing-instance all
Interface    Admin   Link   Proto    Local             Remote Instance
at-0/0/1     up      up     inet     10.0.0.1/24
ge-0/0/0.0   up      up     inet      192.168.4.28/24          sample-a
at-0/1/0.0   up      up     inet6     fe80::a:0:0:4/64         sample-b
so-0/0/0.0   up      up     inet     10.0.0.1/32
```

show interfaces snmp-index (SRX devices)

```
user@host> show interfaces snmp-index 33
Physical interface: so-2/1/1, Enabled, Physical link is Down
  Interface index: 149, SNMP ifIndex: 33
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags    : Present Running Down
  Interface flags: Hardware-Down Point-To-Point SNMP-Traps 16384
  Link flags      : Keepalives
  CoS queues      : 8 supported
  Last flapped    : 2005-06-15 11:45:57 PDT (05:38:43 ago)
  Input rate      : 0 bps (0 pps)
  Output rate     : 0 bps (0 pps)
  SONET alarms    : LOL, PLL, LOS
  SONET defects   : LOL, PLL, LOF, LOS, SEF, AIS-L, AIS-P
```

show interfaces source-class all (SRX devices)

```
user@host> show interfaces source-class all
Logical interface so-0/1/0.0
                                              Packets                 Bytes
      Source class                  (packet-per-second)    (bits-per-second)
                             gold             1928095             161959980
```

 Copyright © 2018, Juniper Networks, Inc.

```
                                 (                889) (             597762)
                          bronze                    0                     0
                                 (                  0) (                  0)
                          silver                    0                     0
                                 (                  0) (                  0)
  Logical interface so-0/1/3.0
                                              Packets                 Bytes
      Source class               (packet-per-second)    (bits-per-second)
                            gold                    0                     0
                                 (                  0) (                  0)
                          bronze                    0                     0
                                 (                  0) (                  0)
                          silver               116113               9753492
                                 (                939) (             631616)
```

show interfaces statistics (Fast Ethernet SRX devices)

```
user@host> show interfaces fe-1/3/1 statistics
Physical interface: fe-1/3/1, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 1042
  Description: ford fe-1/3/1
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:90:69:93:04:dc, Hardware address: 00:90:69:93:04:dc
  Last flapped   : 2006-04-18 03:08:59 PDT (00:01:24 ago)
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 0, Output errors: 0
  Active alarms  : None
  Active defects : None
  Logical interface fe-1/3/1.0 (Index 69) (SNMP ifIndex 50)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500
      Flags: Is-Primary, DCU, SCU-in
                                              Packets                 Bytes
      Destination class          (packet-per-second)    (bits-per-second)
                         silver1                    0                     0
                                 (                  0) (                  0)
                         silver2                    0                     0
                                 (                  0) (                  0)
                         silver3                    0                     0
                                 (                  0) (                  0)
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 10.27.245/24, Local: 10.27.245.2,
        Broadcast: 10.27.245.255
    Protocol iso, MTU: 1497
      Flags: Is-Primary
```

show interfaces switch-port (SRX devices)

```
user@host# show interfaces ge-slot/0/0 switch-port port-number
Port 0, Physical link is Up
        Speed: 100mbps, Auto-negotiation: Enabled
   Statistics:                           Receive          Transmit
          Total bytes                     28437086         21792250
          Total packets                     409145            88008
```

```
          Unicast packets                    9987            83817
          Multicast packets                145002                0
          Broadcast packets                254156             4191
          Multiple collisions                  23               10
          FIFO/CRC/Align errors                 0                0
          MAC pause frames                      0                0
          Oversized frames                      0
          Runt frames                           0
          Jabber frames                         0
          Fragment frames                       0
          Discarded frames                      0
        Autonegotiation information:
          Negotiation status: Complete
          Link partner:
              Link mode: Full-duplex, Flow control: None, Remote fault: OK, Link
 partner Speed: 100 Mbps
          Local resolution:
              Flow control: None, Remote fault: Link OK
```

show interfaces transport pm (SRX devices)

```
user@host> show interfaces transport pm all current et-0/1/0
Physical interface: et-0/1/0, SNMP ifIndex 515
  14:45-current           Elapse time:900 Seconds
  Near End       Suspect Flag:False             Reason:None
    PM                 COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED

  OTU-BBE                0                800              No              No
  OTU-ES                 0                135              No              No
  OTU-SES                0                90               No              No
  OTU-UAS                427              90               No              No
  Far End        Suspect Flag:True              Reason:Unknown
    PM                 COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED

  OTU-BBE                0                800              No              No
  OTU-ES                 0                135              No              No
  OTU-SES                0                90               No              No
  OTU-UAS                0                90               No              No
  Near End       Suspect Flag:False             Reason:None
    PM                 COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED

  ODU-BBE                0                800              No              No
  ODU-ES                 0                135              No              No
  ODU-SES                0                90               No              No
  ODU-UAS                427              90               No              No
  Far End        Suspect Flag:True              Reason:Unknown
    PM                 COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED

  ODU-BBE                0                800              No              No
  ODU-ES                 0                135              No              No
  ODU-SES                0                90               No              No
  ODU-UAS                0                90               No              No
  FEC            Suspect Flag:False             Reason:None
    PM                 COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED

  FEC-CorrectedErr       2008544300       0                NA              NA
  FEC-UncorrectedWords   0                0                NA              NA
  BER            Suspect Flag:False             Reason:None
    PM                MIN       MAX        AVG        THRESHOLD      TCA-ENABLED
     TCA-RAISED
  BER                 3.6e-5    5.8e-5     3.6e-5     10.0e-3        No
```


Copyright © 2018, Juniper Networks, Inc.

```
       Yes
Physical interface: et-0/1/0, SNMP ifIndex 515
  14:45-current
  Suspect Flag:True              Reason:Object Disabled
    PM                            CURRENT   MIN       MAX      AVG       THRESHOLD
      TCA-ENABLED       TCA-RAISED
                                                                            (MIN)
(MAX)    (MIN)  (MAX)        (MIN)  (MAX)
  Lane chromatic dispersion        0        0         0         0         0
 0        NA     NA          NA     NA
  Lane differential group delay    0        0         0         0         0
 0        NA     NA          NA     NA
    q Value                          120      120       120       120       0
   0        NA     NA          NA     NA
  SNR                               28       28        29        28        0
 0        NA     NA          NA     NA
  Tx output power(0.01dBm)         -5000    -5000     -5000     -5000     -300
 -100     No     No          No     No
  Rx input power(0.01dBm)          -3642    -3665     -3626     -3637     -1800
 -500     No     No          No     No
  Module temperature(Celsius)       46       46        46        46        -5
 75       No     No          No     No
  Tx laser bias current(0.1mA)     0        0         0         0         0
 0        NA     NA          NA     NA
  Rx laser bias current(0.1mA)     1270     1270      1270      1270      0
 0        NA     NA          NA     NA
  Carrier frequency offset(MHz)    -186     -186      -186      -186      -5000
 5000     No     No          No     No
```

show security zones (SRX devices)

```
user@host> show security zones
Functional zone: management
   Description: This is the management zone.
   Policy configurable: No
   Interfaces bound: 1
   Interfaces:
     ge-0/0/0.0
Security zone: Host
   Description: This is the host zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     fxp0.0
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
Security zone: def
   Description: This is the def zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/2.0
```

show interfaces (ATM)

Syntax

show interfaces at-*fpc/pic/port*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced before Junos OS Release 7.4.

Description

(M Series and T Series routers only) Display status information about the specified ATM interface.

Options

at-*fpc/pic/port*—Display standard information about the specified ATM interface.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display the SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

List of Sample Output

show interfaces (ATM, IMA Group) on page 1213
show interfaces extensive (ATM IMA Group) on page 1214
show interfaces (ATM1, SONET Mode) on page 1215
show interfaces brief (ATM1, SONET Mode) on page 1216
show interfaces detail (ATM1, SONET Mode) on page 1216
show interfaces extensive (ATM1, SONET Mode) on page 1217
show interfaces (ATM2, SDH Mode) on page 1219
show interfaces brief (ATM2, SDH Mode) on page 1220
show interfaces detail (ATM2, SDH Mode) on page 1221
show interfaces extensive (ATM2, SDH Mode) on page 1222
show interfaces (ATM2, SONET Mode) on page 1225
show interfaces brief (ATM2, SONET Mode) on page 1226
show interfaces detail (ATM2, SONET Mode) on page 1227
show interfaces extensive (ATM2, SONET Mode) on page 1229

Output Fields

Table 59 on page 1199 lists the output fields for the **show interfaces** (ATM) command. Output fields are listed in the approximate order in which they appear.

Copyright © 2018, Juniper Networks, Inc.

Table 59: ATM show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under “Common Output Fields Description” on page 1110.	All levels
Description	Configured interface description.	All levels
Interface index	Physical interface's index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Link-level type	Encapsulation being used on the physical interface: • **ATM-CCC-CELL-RELAY**—ATM cell relay for CCC. • **ATM-CCC-VC-MUX**—ATM virtual circuit (VC) for CCC. • **ATM-CISCO-NLPID**—Cisco-compatible ATM NLPID encapsulation. • **ATM-MIPP-LLC**—ATM MLPPP over ATM Adaptation Layer 5 (AAL5)/logical link control (LLC). • **ATM-NLPID**—ATM NLPID encapsulation. • **ATM-PPP-LLC**—ATM PPP over AAL5/LLC. • **ATM-PPP-VC-MUX**—ATM PPP over raw AAL5. • **ATM-PVC**—ATM permanent virtual circuits. • **ATM-SNAP**—ATM LLC/SNAP encapsulation. • **ATM-TCC-SNAP**—ATM LLC/SNAP for translational cross-connection. • **ATM-TCC-VC-MUX**—ATM VC for translational cross-connection. • **ATM-VC-MUX**—ATM VC multiplexing. • **ETHER-OVER-ATM-LLC**—Ethernet over ATM (LLC/SNAP) encapsulation. • **ETHER-VPLS-OVER-ATM-LLC**—Ethernet VPLS over ATM (bridging) encapsulation.	All levels
MTU	MTU size on the physical interface.	All levels
Clocking	Reference clock source: **Internal** or **External**.	All levels
***framing* Mode**	Framing mode: **SONET** or **SDH**.	All levels
Speed	Speed at which the interface is running as represented by the interface type (for example, **OC3**, **ADSL2+**, and **SHDSL(2-wire)**.	All levels
Loopback	Whether loopback is enabled and the type of loopback (**local** or **remote**).	All levels
Payload scrambler	Whether payload scrambling is enabled.	All levels

Copyright © 2018, Juniper Networks, Inc.

Table 59: ATM show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under "Common Output Fields Description" on page 1110.	All levels
Link flags	Information about the link. Possible values are described in the "Link Flags" section under "Common Output Fields Description" on page 1110.	All levels
CoS queues	Number of CoS queues configured.	**detail extensive** none
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Current address	Ethernet MAC address for this interface for Ethernet over ATM encapsulation.	**detail extensive** none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour:minute:second timezone*** **(*hour:minute:second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.	**detail extensive** none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Statistics for traffic on the interface. • **Input bytes**—Number of bytes received on the interface • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 59: ATM show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface whose definitions are as follows: • **Errors**—Sum of the incoming frame aborts and frame check sequence (FCS) errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's random early detection (RED) mechanism. • **Invalid VCs**—Number of cells that arrived for a nonexistent VC. • **Framing errors**—Sum of AAL5 packets that have FCS errors, reassembly timeout errors, and length errors. • **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. • **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. • **Resource errors**—Sum of transmit drops.	**extensive**
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If it 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. • **Aged packets**—Number of packets that remained so long in shared packet SDRAM 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. • **MTU errors**—Number of packets larger than the MTU threshold. • **Resource errors**—Sum of transmit drops.	**extensive**
Egress queues	Total number of egress queues supported on the specified interface.	**detail extensive**

Table 59: ATM show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Queue counters	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism. NOTE: Physical interface queue counters of ATM2 PICs displayed by the **show interfaces at-*fpc/pic/port* detail** command show the packet forwarding stream statistics associated with the ATM2 ports. Since multiple ports of the ATM2 PICs (except for the ATM2 dual-port OC12) share one packet forwarding stream, the physical interface queue counters reflect the aggregate of ATM2 port statistics.	**detail extensive**
SONET alarms **SONET defects**	SONET media-specific defects that prevent the interface from passing packets. When a defect persists for a certain period, 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 light the red or yellow alarm LED on the craft interface. See these fields for possible alarms and defects: **SONET PHY**, **SONET section**, **SONET line**, and **SONET path**.	**detail extensive** none
SONET PHY	Counts of specific SONET errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **PLL Lock**—Phase-locked loop • **PHY Light**—Loss of optical signal	**extensive**
SONET section	Counts of specific SONET errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **BIP-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**

 Copyright © 2018, Juniper Networks, Inc.

Table 59: ATM show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
SONET line	Active alarms and defects, plus counts of specific SONET errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **BIP-B2**—Bit interleaved parity for SONET line overhead • **REI-L**—Remote error indication (near-end line) • **RDI-L**—Remote defect indication (near-end line) • **AIS-L**—Alarm indication signal (near-end line) • **BERR-SF**—Bit error rate fault signal failure • **BERR-SD**—Bit error rate defect signal degradation • **ES-L**—Errored seconds (near-end line) • **SES-L**—Severely errored seconds (near-end line) • **UAS-L**—Unavailable seconds (near-end line) • **ES-LFE**—Errored seconds (far-end line) • **SES-LFE**—Severely errored seconds (far-end line) • **UAS-LFE**—Unavailable seconds (far-end line)	**extensive**
SONET path	Active alarms and defects, plus counts of specific SONET errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **BIP-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) • **ES-PFE**—Errored seconds (far-end STS path) • **SES-PFE**—Severely errored seconds (far-end STS path) • **UAS-PFE**—Unavailable seconds (far-end STS path)	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 59: ATM show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Received SONET overhead **Transmitted SONET overhead**	Values of the received and transmitted SONET overhead: • **C2**—Signal label. Allocated to identify the construction and content of the STS-level SPE and for PDI-P. • **F1**—Section user channel byte. This byte is set aside for the purposes of users. • **K1** and **K2**—These bytes are allocated for APS signaling for the protection of the multiplex section. • **J0**—Section trace. This byte is defined for STS-1 number 1 of an STS-*N* signal. Used to transmit a 1-byte fixed-length string or a 16-byte message so that a receiving terminal in a section can verify its continued connection to the intended transmitter. • **S1**—Synchronization status. The S1 byte is located in the first STS-1 of an STS-*N*. • **Z3** and **Z4**—Allocated for future use.	**extensive**
SDH alarms **SDH defects**	SDH media-specific defects that can prevent the interface from passing packets. When a defect persists for a certain period, 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 light the red or yellow alarm LED on the craft interface. See these fields for possible alarms and defects: **SDH PHY**, **SDH regenerator section**, **SDH multiplex section**, and **SDH path**.	All levels
SDH PHY	Active alarms and defects, plus counts of specific SDH errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **PLL Lock**—Phase-locked loop • **PHY Light**—Loss of optical signal	**extensive**
SDH regenerator section	Active alarms and defects, plus counts of specific SDH errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **RS-BIP8**—24-bit BIP for multiplex section overhead (B2 bytes) • **OOF**—Out of frame • **LOS**—Loss of signal • **LOF**—Loss of frame • **RS-ES**—Errored seconds (near-end regenerator section) • **RS-SES**—Severely errored seconds (near-end regenerator section) • **RS-SEFS**—Severely errored framing seconds (regenerator section)	**extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 59: ATM show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
SDH multiplex section	Active alarms and defects, plus counts of specific SDH errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **MS-BIP24**—8-bit BIP for high-order path overhead (B3 byte) • **MS-FEBE**—Far-end block error (multiplex section) • **MS-FERF**—Far-end remote fail (multiplex section) • **MS-AIS**—Alarm indication signal (multiplex section) • **BERR-SF**—Bit error rate fault (signal failure) • **BERR-SD**—Bit error rate defect (signal degradation) • **MS-ES**—Errored seconds (near-end multiplex section) • **MS-SES**—Severely errored seconds (near-end multiplex section) • **MS-UAS**—Unavailable seconds (near-end multiplex section) • **MS-ES-FE**—Errored seconds (far-end multiplex section) • **MS-SES-FE**—Severely errored seconds (far-end multiplex section) • **MS-UAS-FE**—Unavailable seconds (far-end multiplex section)	**extensive**
SDH path	Active alarms and defects, plus counts of specific SDH errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **HP-BIP8**—8-bit BIP for regenerator section overhead (B1 byte) • **HP-FEBE**—Far-end block error (high-order path) • **HP-LOP**—Loss of pointer (high-order path) • **HP-AIS**—High-order-path alarm indication signal • **HP-FERF**—Far-end remote fail (high-order path) • **HP-UNEQ**—Unequipped (high-order path) • **HP-PLM**—Payload label mismatch (high-order path) • **HP-ES**—Errored seconds (near-end high-order path) • **HP-SES**—Severely errored seconds (near-end high-order path) • **HP-UAS**—Unavailable seconds (near-end high-order path) • **HP-ES-FE**—Errored seconds (far-end high-order path) • **HP-SES-FE**—Severely errored seconds (far-end high-order path) • **HP-UAS-FE**—Unavailable seconds (far-end high-order path)	**extensive**

Table 59: ATM show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Received SDH overhead **Transmitted SDH overhead**	Values of the received and transmitted SONET overhead: • **C2**—Signal label. This byte is allocated to identify the construction and content of the STS-level SPE and for PDI-P. • **F1**—Section user channel byte. This byte is set aside for the purposes of users. • **K1** and **K2**—These bytes are allocated for APS signaling for the protection of the multiplex section. • **J0**—Section trace. This byte is defined for STS-1 number 1 of an STS-*N* signal. This bye is used to transmit a 1-byte fixed-length string or a 16-byte message so that a receiving terminal in a section can verify its continued connection to the intended transmitter. • **S1**—Synchronization status. The S1 byte is located in the first STS-1 of an STS-*N*. • **Z3** and **Z4**—These bytes are allocated for future use.	**extensive**
Received path trace **Transmitted path trace**	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**
ATM Status	ATM state information: • **HCS State**—Status of the header check sequence. ATM uses the HCS field in the cell header in the cell delineation process to frame ATM cell boundaries. The HCS is an FCS-8 calculation over the first four octets of the ATM cell header. • **LOC**—Current loss of cell (LOC) delineation state. **OK** means that no LOC is currently asserted.	**extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 59: ATM show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
ATM Statistics	ATM statistics for the interface: • **Uncorrectable HCS errors**—Number of cells dropped because the cell delineation failed. These errors most likely indicate that a SONET/SDH layer problem has occurred. • **Correctable HCS errors**—Number of correctable HCS errors that occurred. The cell delineation process can recover from these errors and locate the ATM cell boundary, although the framing process is not quite stable. The ATM cell is not dropped. This counter increases when the cell delineation process changes its state from **present** to **sync** (for example, when a cable is plugged into the interface). The following error statistics are from the framer: • **Tx cell FIFO overruns**—Number of overruns in the transmit FIFO. • **Rx cell FIFO overruns**—Number of overruns in the receive FIFO. • **Rx cell FIFO underruns**—Number of underruns in the receive FIFO. • **Input cell count**—Number of ATM cells received by the interface (not including idle cells). • **Output cell count**—Number of ATM cells transmitted by the interface (including idle cells). • **Output idle cell count**—Number of idle cells sent by the port. When ATM has nothing to send, it sends idle cells to fill the time slot. • **Output VC queue drops**—Number of packets dropped by a port on the PIC. Packets are dropped because of queue limits on the VCs. The following error statistics are from the SAR: • **Input no buffers**—Number of AAL5 packets dropped because no channel blocks or buffers were available to handle them. • **Input length errors**—Number of AAL5 packets dropped because their length was incorrect. Usually, these errors occur because a cell has been corrupted or lost, or because the length field was corrupted. They can also mean the AAL5 length field was zero. • **Input timeouts**—Number of AAL5 packets dropped because of a reassembly timeout. • **Input invalid VCs**—Number of AAL5 packets dropped because the header was unrecognized (because the VC was not correct or not configured). • **Input bad CRCs**—Number of AAL5 packets dropped because of frame check sequence errors. • **Input OAM cell no buffers**—Number of received OAM cells or raw cells dropped because no buffers were available to handle them. • **L2 circuit out-of-sequence packets**—(Layer 2 AAL5 mode) Number of AAL5 packets that are out of sequential order. • **Denied packets count**—The number of packets dropped due to VLAN priority deny packets or due to an error forwarding configuration that might cause a negative frame length, that is, the stripping size is larger than the packet size.	**extensive**
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number.	**extensive**

Table 59: ATM show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface. • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 59: ATM show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output

VPI	(ATM2) Virtual path identifier information:
	• **Flags**—VPI flags can be one or more of the following:
	  ◦ **Active** (virtual path is up)
	  ◦ **OAM** (operation and maintenance is enabled)
	  ◦ **Shaping** (shaping is configured)
	• **CBR, Peak**
	• **OAM, Period**—Interval at which OAM F4 loopback cells are sent.
	• **Up count**—Number of F4 OAM cells required to consider the virtual path up; the range is **1** through **255**.
	• **Down count**—Number of F4 OAM cells required to consider the virtual path down; the range is **1** through **255**.
	• **Total down time**—Total number of seconds the VPI has been down since it was opened, using the format **Total down time:** ***hh:mm:ss*** or **Never.**
	• **Last down**—Time of last **Down** transition, using the format **Last down:** ***hh:mm:ss*** **ago** or **Never**.
	• **OAM F4 cell statistics**—(Nonpromiscuous mode) OAM F4 statistics:
	  ◦ **Total received**—Number of OAM F4 cells received.
	  ◦ **Total sent**—Number of OAM F4 cells sent.
	  ◦ **Loopback received**—Number of OAM F4 loopback cells received.
	  ◦ **Loopback sent**—Number of OAM F4 loopback cells sent.
	  ◦ **Last received**—Time at which the last OAM F4 cell was received.
	  ◦ **Last sent**—Time at which the last OAM F4 cell was sent.
	  ◦ **RDI received**—Number of OAM F4 cells received with the remote defect indication bit set.
	  ◦ **RDI sent**—Number of OAM F4 cells sent with the RDI bit set.
	  ◦ **AIS received**—Number of OAM F4 cells received with the alarm indication signal bit set.
	  ◦ **AIS sent**—Number of OAM F4 cells sent with the AIS bit set.
	Traffic statistics:
	• **Input bytes**—Number of bytes received on the VPI.
	• **Output bytes**—Number of bytes transmitted on the VPI.
	• **Input packets**—Number of packets received on the VPI.
	• **Output packets**—Number of packets transmitted on the VPI.
		detail extensive none

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**

Copyright © 2018, Juniper Networks, Inc.

Table 59: ATM 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 1110.	All levels
Input packets	Number of packets received on the logical interface.	None specified
Output packets	Number of packets transmitted on the logical interface.	None specified
Encapsulation	Encapsulation on the logical interface.	All levels
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.	**detail extensive**
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.	**detail extensive**
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.	**detail extensive**
Input packets	Number of packets received on the logical interface.	None specified
Output packets	Number of packets transmitted on the logical interface.	None specified
protocol-family	Protocol family configured on the logical interface. If the protocol is **inet**, the IP address of the interface is also displayed.	**brief**
Protocol	Protocol family configured on the logical interface.	**detail extensive** none
MTU	MTU size on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route table	Routing table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive**
Flags	Information about the protocol family flags. Possible values are described in the "Family Flags" section under "Common Output Fields Description" on page 1110.	**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 1110.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none

 Copyright © 2018, Juniper Networks, Inc.

Table 59: ATM show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

VCI	Virtual circuit identifier number and information:
	• **Flags**—VCI flags:
	  ◦ **Active**—VCI is up and in working condition.
	  ◦ **CCC down**—VCI CCC is not in working condition.
	  ◦ **Closed**—VCI is closed because the user disabled the logical or physical interface from the CLI.
	  ◦ **Configured**—VCI is configured.
	  ◦ **Down**—VCI is not in working condition. The VCI might have alarms, defects, F5 AIS/RDI, or no response to OAM loopback cells.
	  ◦ **ILMI**—VCI is up and in working condition.
	  ◦ **OAM**—OAM loopback is enabled.
	  ◦ **Multicast**—VCI is a multicast VCI or DLCI.
	  ◦ **Multipoint destination**—VCI is configured as a multipoint destination.
	  ◦ **None**—No VCI flags.
	  ◦ **Passive-OAM**—Passive OAM is enabled.
	  ◦ **Shaping**—Shaping is enabled.
	  ◦ **Sustained**—Shaping rate is set to **Sustained**.
	  ◦ **Unconfigured**—VCI is not configured.
	• **Total down time**—Total number of seconds the VCI has been down, using the format **Total down time:** ***hh*:*mm*:*ss*** or**Never**.
	• **Last down**—Time of last **Down** transition, using the format **Last down:** ***hh*:*mm*:*ss***.
	• **EPD threshold**—(ATM2 only) Threshold at which a packet is dropped when the queue size (in number of cells) exceeds the early packet-discard (EPD) value.
		All levels

Table 59: ATM show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output

VCI (continued)	• **Transmit weight cells**—(ATM2 only) Amount of bandwidth assigned to this queue.
	• **ATM per-VC transmit statistics**:
	  ◦ **Tail queue packet drops**—Number of packets dropped because of bandwidth constraints. This value indicates that packets are queued to send out at a rate faster than allowed.
	• **OAM F4 cell statistics**—(Nonpromiscuous mode) OAM F4 statistics:
	  ◦ **Total received**—Number of OAM F4 cells received.
	  ◦ **Total sent**—Number of OAM F4 cells sent.
	  ◦ **Loopback received**—Number of OAM F4 loopback cells received.
	  ◦ **Loopback sent**—Number of OAM F4 loopback cells sent.
	  ◦ **Last received**—Time at which the last OAM F4 cell was received.
	  ◦ **Last sent**—Time at which the last OAM F4 cell was sent.
	  ◦ **RDI received**—Number of OAM F4 cells received with the remote defect indication bit set.
	  ◦ **RDI sent**—Number of OAM F4 cells sent with the RDI bit set.
	  ◦ **AIS received**—Number of OAM F4 cells received with the alarm indication signal bit set.
	  ◦ **AIS sent**—Number of OAM F4 cells sent with the AIS bit set.
	• **Traffic statistics**—Number and rate of bytes and packets received and transmitted on the physical interface.
	  ◦ **Input bytes**—Number of bytes received on the interface.
	  ◦ **Output bytes**—Number of bytes transmitted on the interface.
	  ◦ **Input packets**—Number of packets received on the interface
	  ◦ **Output packets**—Number of packets transmitted on the interface.
		All levels

IMA group properties	• **Version**—The specified IMA specification version, either IMA 1.0 or IMA 1.1.
	• **Frame length**—The specified frame size, which can be 32, 64, 128, or 256.
	• **Differential delay**—Maximum differential delay among links in milliseconds.
	• **Symmetry**—Either Common Transmit Clock or Independent Transmit Clock timing mode.
	• **Transmit clock**—The specified IMA clock mode, either common or independent.
	• **Minimum links**—The number of minimum active links specified in both transmit and receive directions.
	  ◦ **Transmit**—The per-PIC limit on the number of minimum active links in the transmit direction.
	  ◦ **Receive**—The per-PIC limit on the number of minimum active links in the receive direction.
	• **Frame synchronization**—The specified IMA frame synchronization state transition variables (Alpha, Beta, and Gamma) and their specified values.
	  ◦ **Alpha**—The number of consecutive invalid ICP cells for IFSM.
	  ◦ **Beta**—The number of consecutive errored ICP cells for IFSM.
	  ◦ **Gamma**—The number of consecutive valid ICP cells for IFSM.
	• **Links**—The number of IMA links assigned to the IMA group.
		detail extensive none

Copyright © 2018, Juniper Networks, Inc.

Table 59: ATM show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
IMA group alarms	• **Start-up-FE**—Far-end group alarm status • **Config-Aborted**—Near-end configuration aborted group alarm status • **Config-Aborted-FE**—Far-end configuration aborted group alarm status • **Insufficient-Links**—Near-end insufficient links group alarm status • **Insufficient-Links-FE**—Far-end insufficient links group alarm status • **Blocked-FE**—Far-end blocked group alarm status • **GR-Timing-Mismatch**—Group timing mismatch alarm status	**detail extensive** none
IMA group defects	• **Start-up-FE**—Far-end group defect status • **Config-Aborted**—Near-end configuration aborted group defect status • **Config-Aborted-FE**—Far-end configuration aborted group defect status • **Insufficient-Links**—Near-end insufficient links group defect status • **Insufficient-Links-FE**—Far-end insufficient links group defect status • **Blocked-FE**—Far-end blocked group defect status • **GR-Timing-Mismatch**—Group timing mismatch defect status	**detail extensive** none
IMA Group state	Near-end and far-end group status	**detail extensive** none
IMA group media	IMA group media status, including seconds, count and state for the following media parameters: • **FC** • **FC-FE** • **Addr-Mismatch** • **Running** • **UAS**	**detail extensive** none

Sample Output

show interfaces (ATM, IMA Group)

```
user@host> show interfaces at-1/0/0
Physical interface: at-1/0/0, Enabled, Physical link is Up
  IMA group properties:
    Version              : 1.1
    Frame length         : 128
    Differential delay   : 25 milliseconds
    Symmetry             : Symmetrical Configuration and Operation
    Transmit clock       : Common
    Minimum links        : Transmit: 1, Receive: 1
    Frame synchronization: Alpha: 2, Beta: 2, Gamma: 1
    Links                : None
  IMA group alarms   : Start-up-FE Config-Aborted Config-Aborted-FE
Insufficient-Links Insufficient-Links-FE Blocked-FE GR-Timing-Mismatch
  IMA group defects : Start-up-FE Config-Aborted Config-Aborted-FE
Insufficient-Links Insufficient-Links-FE Blocked-FE GR-Timing-Mismatch
  IMA Group state:
    Near end : Start up
    Far  end : Start up
```

Copyright © 2018, Juniper Networks, Inc.

```
IMA group media:         Seconds         Count  State
  FC                                         0
  FC-FE                                      0
  Addr-Mismatch                              0
  Running                      0
  UAS                          0
```

show interfaces extensive (ATM IMA Group)

```
user@host> show interfaces at-0/0/10 extensive
Physical interface: at-0/0/10, Enabled, Physical link is Up
   Interface index: 178, SNMP ifIndex: 540, Generation: 531
  Link-level type: ATM-PVC, MTU: 2048, Speed: Unspecified, Loopback: None, Payload
 scrambler: Enabled
   Device flags   : Present Running
   Link flags     : None
   CoS queues     : 8 supported, 4 maximum usable queues
   Hold-times     : Up 0 ms, Down 0 ms
   Current address: 00:00:5e:00:53:0a
   Last flapped   : 2012-03-16 16:49:15 PDT (2d 07:12 ago)
   Statistics last cleared: 2012-03-16 16:56:58 PDT (2d 07:05 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, Invalid VCs: 0, Framing errors: 0, Policed discards:
0, L3 incompletes: 0, L2 channel errors: 0,
     L2 mismatch timeouts: 0, Resource errors: 0
   Output errors:
     Carrier transitions: 0, Errors: 0, Drops: 0, Aged packets: 0, MTU errors:
0, Resource errors: 0
   IMA group properties:
     Version              : 1.1
     Frame length         : 128
     Differential delay   : 25 milliseconds
     Symmetry             : Symmetrical Configuration and Operation
     Transmit clock       : Common
     Minimum links        : Transmit: 1, Receive: 1
     Frame synchronization: Alpha: 2, Beta: 2, Gamma: 1
     Link #1              : t1-0/0/4             up
   IMA Group alarms    : None
   IMA Group defects   : None

   IMA Group state:
     Near end : Operational
     Far  end : Operational
   IMA group media:         Seconds         Count  State
     FC                                         0
     FC-FE                                      0
     Addr-Mismatch                              0
     Running                 198306
     UAS                          0
   ATM status:
     HCS state:     Sync
```


Copyright © 2018, Juniper Networks, Inc.

```
    LOC      :       OK
  ATM Statistics:
    Uncorrectable HCS errors: 0, Correctable HCS errors: 0, Tx cell FIFO overruns:
 0, Rx cell FIFO overruns: 0,
    Rx cell FIFO underruns: 0, Input cell count: 0, Output cell count: 0, Output
 idle cell count: 0,
    Output VC queue drops: 0, Input no buffers: 0, Input length errors: 0, Input
 timeouts: 0, Input invalid VCs: 0,
    Input bad CRCs: 0, Input OAM cell no buffers: 0
  Packet Forwarding Engine configuration:
    Destination slot: 0
    VPI 2
      Flags: Active
      Total down time: 0 sec, Last down: Never
        Traffic statistics:
         Input  bytes  :                    0
         Output bytes  :                    0
         Input  packets:                    0
         Output packets:                    0

   Logical interface at-0/0/10.602 (Index 71) (SNMP ifIndex 1057) (Generation
17226)
     Flags: Point-To-Point SNMP-Traps CCC-Down 0x0 Encapsulation:
ATM-CCC-Cell-Relay
     L2 circuit cell bundle size: 1, bundle timeout: 125 usec, timeout count: 0
     L2 circuit out-of-sequence count: 0, denied packets count: 0
```

show interfaces (ATM1, SONET Mode)

```
user@host> show interfaces at-1/0/0
Physical interface: at-1/0/0, Enabled, Physical link is Up
  Interface index: 300, SNMP ifIndex: 194
  Description: to allspice at-1/0/0
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SONET mode,
  Speed: OC3, Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:00:5e:00:53:fe
  Last flapped   : 2006-02-24 14:28:12 PST (6d 01:51 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : None
  SONET defects  : None

  Logical interface at-1/0/0.0 (Index 64) (SNMP ifIndex 204)
    Flags: Point-To-Point SNMP-Traps Encapsulation: ATM-SNAP
    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: 4470
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.220.24/30, Local: 192.168.220.26,
        Broadcast: 192.168.220.27
    Protocol iso, MTU: 4470
      Flags: None
    VCI 0.128
      Flags: Active
      Total down time: 0 sec, Last down: Never
        Input packets : 0
        Output packets: 0
```

show interfaces brief (ATM1, SONET Mode)

```
user@host> show interfaces at-1/0/0 brief
Physical interface: at-1/0/0, Enabled, Physical link is Up
  Description: to allspice at-1/0/0
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SONET mode,
  Speed: OC3, Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None

  Logical interface at-1/0/0.0
    Flags: Point-To-Point SNMP-Traps Encapsulation: ATM-SNAP
    inet  192.168.220.26/30
    iso
    VCI 0.128
      Flags: Active
      Total down time: 0 sec, Last down: Never
```

show interfaces detail (ATM1, SONET Mode)

```
user@host> show interfaces at-1/0/0 detail
Physical interface: at-1/0/0, Enabled, Physical link is Up
  Interface index: 300, SNMP ifIndex: 194, Generation: 183
  Description: to allspice at-1/0/0
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SONET mode,
  Speed: OC3, Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:fe
  Last flapped   : 2006-02-24 14:28:12 PST (6d 01:55 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
  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

  SONET alarms   : None
  SONET defects  : None

  Logical interface at-1/0/0.0 (Index 64) (SNMP ifIndex 204) (Generation 5)
    Flags: Point-To-Point SNMP-Traps Encapsulation: ATM-SNAP
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Local statistics:
```

 Copyright © 2018, Juniper Networks, Inc.

```
     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: 4470, Generation: 13, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.220.24/30, Local: 192.168.220.26,
        Broadcast: 192.168.220.27, Generation: 14
    Protocol iso, MTU: 4470, Generation: 14, Route table: 0
      Flags: None
    VCI 0.128
      Flags: Active
      Total down time: 0 sec, Last down: Never
      ATM per-VC transmit statistics:
       Tail queue packet drops: 0
      Traffic statistics:
       Input  bytes  :                    0
       Output bytes  :                    0
       Input  packets:                    0
       Output packets:                    0
```

show interfaces extensive (ATM1, SONET Mode)

```
user@host> show interfaces at-1/0/0 extensive
Physical interface: at-1/0/0, Enabled, Physical link is Up
  Interface index: 300, SNMP ifIndex: 194, Generation: 183
  Description: to allspice at-1/0/0
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SONET mode,
  Speed: OC3, Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:fe
  Last flapped   : 2006-02-24 14:28:12 PST (6d 01:56 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, Invalid VCs: 0, Framing errors: 0, Policed discards: 0,

    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    Resource errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, Aged packets: 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                    0                    0                    0
```

Copyright © 2018, Juniper Networks, Inc.

```
  1 expedited-fo                   0                    0                    0

  2 assured-forw                   0                    0                    0

  3 network-cont                   0                    0                    0

SONET alarms   : None
SONET defects  : None
SONET PHY:           Seconds        Count  State
  PLL Lock                 0            0  OK
  PHY Light                0            0  OK
SONET section:
  BIP-B1                   0            0
  SEF                      0            0  OK
  LOS                      0            0  OK
  LOF                      0            0  OK
  ES-S                     0
  SES-S                    0
  SEFS-S                   0
SONET line:
  BIP-B2                   0            0
  REI-L                    0            0
  RDI-L                    0            0  OK
  AIS-L                    0            0  OK
  BERR-SF                  0            0  OK
  BERR-SD                  0            0  OK
  ES-L                     0
  SES-L                    0
  UAS-L                    0
  ES-LFE                   0
  SES-LFE                  0
  UAS-LFE                  0
SONET path:
  BIP-B3                   0            0
  REI-P                    0            0
  LOP-P                    0            0  OK
  AIS-P                    0            0  OK
  RDI-P                    0            0  OK
  UNEQ-P                   1            1  OK
  PLM-P                    0            0  OK
  ES-P                     1
  SES-P                    1
  UAS-P                    0
  ES-PFE                   0
  SES-PFE                  0
  UAS-PFE                  0
Received SONET overhead:
  F1      : 0x00, J0      : 0x00, K1      : 0x00, K2      : 0x00
  S1      : 0x00, C2      : 0x13, C2(cmp) : 0x13, F2      : 0x00
  Z3      : 0x00, Z4      : 0x00, S1(cmp) : 0x00
Transmitted SONET overhead:
  F1      : 0x00, J0      : 0x01, K1      : 0x00, K2      : 0x00
  S1      : 0x00, C2      : 0x13, F2      : 0x00, Z3      : 0x00
  Z4      : 0x00
ATM status:
  HCS state:     Sync
  LOC      :       OK
ATM Statistics:
  Uncorrectable HCS errors: 0, Correctable HCS errors: 0,
  Tx cell FIFO overruns: 0, Rx cell FIFO overruns: 0,
  Rx cell FIFO underruns: 0, Input cell count: 0, Output cell count: 0,
```

 Copyright © 2018, Juniper Networks, Inc.

```
    Output idle cell count: 0, Output VC queue drops: 0, Input no buffers: 0,
    Input length errors: 0, Input timeouts: 0, Input invalid VCs: 0,
    Input bad CRCs: 0, Input OAM cell no buffers: 0
  Packet Forwarding Engine configuration:
    Destination slot: 1
  CoS information:
    CoS transmit queue          Bandwidth          Buffer      Priority    Limit
                          %           bps     %      usec
    0 best-effort        95     147744000    95         0          low     none
    3 network-control     5       7776000     5         0          low     none

  Logical interface at-1/0/0.0 (Index 64) (SNMP ifIndex 204) (Generation 5)
    Flags: Point-To-Point SNMP-Traps Encapsulation: ATM-SNAP
    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: 4470, Generation: 13, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.220.24/30, Local: 192.168.220.26,
        Broadcast: 192.168.220.27, Generation: 14
    Protocol iso, MTU: 4470, Generation: 14, Route table: 0
      Flags: None
    VCI 0.128
      Flags: Active
      Total down time: 0 sec, Last down: Never
      ATM per-VC transmit statistics:
       Tail queue packet drops: 0
      Traffic statistics:
       Input  bytes  :                    0
       Output bytes  :                    0
       Input  packets:                    0
       Output packets:                    0
```

show interfaces (ATM2, SDH Mode)

```
user@host> show interfaces at-0/2/1
Physical interface: at-0/2/1, Enabled, Physical link is Up
  Interface index: 154, SNMP ifIndex: 42
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SDH mode, Speed: OC3,

  Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:00:5e:00:53:3f
  Last flapped   : 2006-03-24 13:29:58 PST (00:04:48 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
```

```
SDH   alarms   : None
SDH   defects  : None
  VPI 0
    Flags: Active
    Total down time: 0 sec, Last down: Never
Traffic statistics:
       Input  packets:                    0
       Output packets:                    0

Logical interface at-0/2/1.0 (Index 75) (SNMP ifIndex 51)
  Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: ATM-SNAP
  Input packets : 0
  Output packets: 0
  Protocol inet, MTU: 4470
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 10.0.12.6, Local: 10.0.12.5
  Protocol iso, MTU: 4470
    Flags: None
  VCI 0.128
    Flags: Active
    Total down time: 0 sec, Last down: Never
    EPD threshold: 2129, Transmit weight cells: 0
      Input packets : 0
      Output packets: 0

Logical interface at-0/2/1.32767 (Index 76) (SNMP ifIndex 50)
  Flags: Point-To-Multipoint No-Multicast SNMP-Traps 0x4000
  Encapsulation: ATM-VCMUX
  Input packets : 0
  Output packets: 0
  VCI 0.4
    Flags: Active
    Total down time: 0 sec, Last down: Never
    EPD threshold: 0, Transmit weight cells: 0
      Input packets : 0
      Output packets: 0
```

show interfaces brief (ATM2, SDH Mode)

```
user@host> show interfaces at-0/2/1 brief
Physical interface: at-0/2/1, Enabled, Physical link is Up
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SDH mode,
  Speed: OC3, Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None
  Logical interface at-0/2/1.0
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: ATM-SNAP
    inet  10.0.12.5        --> 10.0.12.6
    iso
    VCI 0.128
      Flags: Active
      Total down time: 0 sec, Last down: Never
      EPD threshold: 2129, Transmit weight cells: 0

  Logical interface at-0/2/1.32767
    Flags: Point-To-Multipoint No-Multicast SNMP-Traps 0x4000
    Encapsulation: ATM-VCMUX
    VCI 0.4
      Flags: Active
```

 Copyright © 2018, Juniper Networks, Inc.

```
      Total down time: 0 sec, Last down: Never
      EPD threshold: 0, Transmit weight cells: 0
```

show interfaces detail (ATM2, SDH Mode)

```
user@host> show interfaces at-0/2/1 detail
Physical interface: at-0/2/1, Enabled, Physical link is Up
  Interface index: 154, SNMP ifIndex: 42, Generation: 40
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SDH mode, Speed: OC3,

  Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:3f
  Last flapped   : 2006-03-24 13:29:58 PST (00:05:10 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
  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

  SDH   alarms   : None
  SDH   defects  : None
    VPI 0
      Flags: Active
      Total down time: 0 sec, Last down: Never
        Traffic statistics:
         Input  bytes  :                    0
         Output bytes  :                    0
         Input  packets:                    0
         Output packets:                    0

  Logical interface at-0/2/1.0 (Index 75) (SNMP ifIndex 51) (Generation 25)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: ATM-SNAP
    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: 4470, Generation: 62, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.0.12.6, Local: 10.0.12.5, Broadcast: Unspecified,
        Generation: 58
    Protocol iso, MTU: 4470, Generation: 63, Route table: 0
      Flags: None
    VCI 0.128
      Flags: Active
      Total down time: 0 sec, Last down: Never
      EPD threshold: 2129, Transmit weight cells: 0
      ATM per-VC transmit statistics:
       Tail queue packet drops: 0
      Traffic statistics:
       Input  bytes  :                    0
       Output bytes  :                    0
       Input  packets:                    0
       Output packets:                    0
  Logical interface at-0/2/1.32767 (Index 76) (SNMP ifIndex 50) (Generation 26)
    Flags: Point-To-Multipoint No-Multicast SNMP-Traps 0x4000
    Encapsulation: ATM-VCMUX
    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
    VCI 0.4
      Flags: Active
      Total down time: 0 sec, Last down: Never
      EPD threshold: 0, Transmit weight cells: 0
      ATM per-VC transmit statistics:
       Tail queue packet drops: 0
      Traffic statistics:
       Input  bytes  :                    0
       Output bytes  :                    0
       Input  packets:                    0
       Output packets:                    0
```

show interfaces extensive (ATM2, SDH Mode)

```
user@host> show interfaces at-0/2/1 extensive
Physical interface: at-0/2/1, Enabled, Physical link is Up
  Interface index: 154, SNMP ifIndex: 42, Generation: 40
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SDH mode, Speed: OC3,

  Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:3f
  Last flapped   : 2006-03-24 13:29:58 PST (00:06:49 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
```

 Copyright © 2018, Juniper Networks, Inc.

```
 Output bytes  :                   0                    0 bps
 Input  packets:                   0                    0 pps
 Output packets:                   0                    0 pps
Input errors:
  Errors: 0, Drops: 0, Invalid VCs: 0, Framing errors: 0, Policed discards: 0,

  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  Resource errors: 0
Output errors:
  Carrier transitions: 3, Errors: 0, Drops: 0, Aged packets: 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                  0                    0                    0

  1 expedited-fo                 0                    0                    0

  2 assured-forw                 0                    0                    0

  3 network-cont                 0                    0                    0

SDH   alarms   : None
SDH   defects  : None
SDH PHY:              Seconds        Count  State
  PLL Lock                  0            0  OK
  PHY Light                 1            1  OK
SDH regenerator section:
  RS-BIP8                   2         8828
  OOF                       2            2  OK
  LOS                       2            1  OK
  LOF                       2            1  OK
  RS-ES                     4
  RS-SES                    3
  RS-SEFS                   2
SDH multiplex section:
  MS-BIP24                  2          771
  MS-FEBE                   1        17476
  MS-FERF                   2            1  OK
  MS-AIS                    2            1  OK
  BERR-SF                   0            0  OK
  BERR-SD                   0            0  OK
  MS-ES                     4
  MS-SES                    2
  MS-UAS                    0
  MS-ES-FE                  3
  MS-SES-FE                 2
  MS-UAS-FE                 0
SDH path:
  HP-BIP8                   1            6
  HP-FEBE                   1          251
  HP-LOP                    0            0  OK
  HP-AIS                    2            1  OK
  HP-FERF                   3            2  OK
  HP-UNEQ                   1            1  OK
  HP-PLM                    2            1  OK
  HP-ES                     4
  HP-SES                    3
  HP-UAS                    0
  HP-ES-FE                  3
```

```
  HP-SES-FE                 3
  HP-UAS-FE                 0
Received SDH overhead:
  F1      : 0x00, J0      : 0x00, K1      : 0x00, K2      : 0x00
  S1      : 0x00, C2      : 0x13, C2(cmp) : 0x13, F2      : 0x00
  Z3      : 0x00, Z4      : 0x00, S1(cmp) : 0x00
Transmitted SDH overhead:
  F1      : 0x00, J0      : 0x01, K1      : 0x00, K2      : 0x00
  S1      : 0x00, C2      : 0x13, F2      : 0x00, Z3      : 0x00
  Z4      : 0x00
ATM status:
  HCS state:     Sync
  LOC      :       OK
ATM Statistics:
  Uncorrectable HCS errors: 0, Correctable HCS errors: 0,
  Tx cell FIFO overruns: 0, Rx cell FIFO overruns: 0,
  Rx cell FIFO underruns: 0, Input cell count: 0, Output cell count: 0,
  Output idle cell count: 0, Output VC queue drops: 0, Input no buffers: 0,
  Input length errors: 0, Input timeouts: 0, Input invalid VCs: 0,
  Input bad CRCs: 0, Input OAM cell no buffers: 0
Packet Forwarding Engine configuration:
  Destination slot: 0
  VPI 0
    Flags: Active
    Total down time: 0 sec, Last down: Never
      Traffic statistics:
       Input  bytes  :                  0
       Output bytes  :                  0
       Input  packets:                  0
       Output packets:                  0

Logical interface at-0/2/1.0 (Index 75) (SNMP ifIndex 51) (Generation 25)
  Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: ATM-SNAP
  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: 4470, Generation: 62, Route table: 0
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 10.0.12.6, Local: 10.0.12.5, Broadcast: Unspecified,
      Generation: 58
  Protocol iso, MTU: 4470, Generation: 63, Route table: 0
    Flags: None
  VCI 0.128
    Flags: Active
    Total down time: 0 sec, Last down: Never
    EPD threshold: 2129, Transmit weight cells: 0
    ATM per-VC transmit statistics:
     Tail queue packet drops: 0
```

 Copyright © 2018, Juniper Networks, Inc.

```
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
Logical interface at-0/2/1.32767 (Index 76) (SNMP ifIndex 50) (Generation 26)
  Flags: Point-To-Multipoint No-Multicast SNMP-Traps 0x4000
  Encapsulation: ATM-VCMUX
  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
  VCI 0.4
    Flags: Active
    Total down time: 0 sec, Last down: Never
    EPD threshold: 0, Transmit weight cells: 0
    ATM per-VC transmit statistics:
     Tail queue packet drops: 0
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
```

show interfaces (ATM2, SONET Mode)

```
user@host> show interfaces at-0/3/1
Physical interface: at-0/3/1, Enabled, Physical link is Up
  Interface index: 139, SNMP ifIndex: 67
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SONET mode,
  Speed: OC3, Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:00:5e:00:53:5e
  Last flapped   : 2006-03-13 17:46:36 PST (16:01:12 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : None
  SONET defects  : None
    VPI 0
      Flags: Active, OAM, Shaping
      CBR, Peak: 50kbps
      OAM, Period 30 sec, Up count: 10, Down count: 10
      Total down time: 0 sec, Last down: Never
     OAM F4 cell statistics:
      Total received: 4, Total sent: 4
      Loopback received: 4, Loopback sent: 4
      RDI received: 0, RDI sent: 0
      AIS received: 0
  Traffic statistics:
         Input  packets:                   4
         Output packets:                  30
    VPI 10
      Flags: Active
```

```
    Total down time: 0 sec, Last down: Never
  Traffic statistics:
       Input  packets:                    0
       Output packets:                    0
  Logical interface at-0/3/1.0 (Index 78) (SNMP ifIndex 77)
    Flags: Point-To-Point Copy-PLP-To-CLP SNMP-Traps 0x4000
    Encapsulation: ATM-SNAP
    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: 4470
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.0.59.5, Local: 10.0.59.6
    Protocol iso, MTU: 4470
      Flags: None
    VCI 0.128
      Flags: Active
      Total down time: 0 sec, Last down: Never
      EPD threshold: 2129, Transmit weight cells: 10
        Input packets : 0
        Output packets: 0

  Logical interface at-0/3/1.32767 (Index 79) (SNMP ifIndex 76)
    Flags: Point-To-Multipoint Copy-PLP-To-CLP No-Multicast SNMP-Traps 0x4000
    Encapsulation: ATM-VCMUX
    Input packets : 4
    Output packets: 30
    VCI 0.16
      Flags: Active, ILMI
      Total down time: 0 sec, Last down: Never
      EPD threshold: 0, Transmit weight cells: 0
        Input packets : 0
        Output packets: 26
    VCI 0.4
      Flags: Active, OAM
      OAM, Period 30 sec, Up count: 10, Down count: 10
      Total down time: 0 sec, Last down: Never
      EPD threshold: 2129, Transmit weight cells: 0
        Input packets : 4
        Output packets: 4
     OAM F4 cell statistics:
      Total received: 4, Total sent: 4
      Loopback received: 4, Loopback sent: 4
      RDI received: 0, RDI sent: 0
      AIS received: 0, AIS sent: 0
```

show interfaces brief (ATM2, SONET Mode)

```
user@host> show interfaces at-0/3/1 brief
Physical interface: at-0/3/1, Enabled, Physical link is Up
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SONET mode,
  Speed: OC3, Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None

  Logical interface at-0/3/1.0
    Flags: Point-To-Point Copy-PLP-To-CLP SNMP-Traps 0x4000
    Encapsulation: ATM-SNAP
    inet  10.0.59.6        --> 10.0.59.5
    iso
    VCI 0.128
```

 Copyright © 2018, Juniper Networks, Inc.

```
    Flags: Active
    Total down time: 0 sec, Last down: Never
    EPD threshold: 2129, Transmit weight cells: 10

Logical interface at-0/3/1.32767
  Flags: Point-To-Multipoint Copy-PLP-To-CLP No-Multicast SNMP-Traps 0x4000
  Encapsulation: ATM-VCMUX
  VCI 0.16
    Flags: Active, ILMI
    Total down time: 0 sec, Last down: Never
    EPD threshold: 0, Transmit weight cells: 0
  VCI 0.4
    Flags: Active, OAM
    Total down time: 0 sec, Last down: Never
    EPD threshold: 2129, Transmit weight cells: 0
```

show interfaces detail (ATM2, SONET Mode)

```
user@host> show interfaces at-0/3/1 detail
Physical interface: at-0/3/1, Enabled, Physical link is Up
  Interface index: 139, SNMP ifIndex: 67, Generation: 22
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SONET mode,
  Speed: OC3, Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:5e
  Last flapped   : 2006-03-13 17:46:36 PST (16:02:39 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                  312                    0 bps
   Output bytes  :                 2952                    0 bps
   Input  packets:                    6                    0 pps
   Output packets:                   50                    0 pps
  Egress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                   44                   44                    0

    1 expedited-fo                   0                    0                    0

    2 assured-forw                   0                    0                    0

    3 network-cont                   6                    6                    0

  SONET alarms   : None
  SONET defects  : None
    VPI 0
      Flags: Active, OAM, Shaping
      CBR, Peak: 50kbps
      OAM, Period 30 sec, Up count: 10, Down count: 10
      Total down time: 0 sec, Last down: Never
     OAM F4 cell statistics:
      Total received: 6, Total sent: 6
      Loopback received: 6, Loopback sent: 6
      Last received: 00:00:29, Last sent: 00:00:29
      RDI received: 0, RDI sent: 0
      AIS received: 0
        Traffic statistics:
         Input  bytes  :                  312
```

Copyright © 2018, Juniper Networks, Inc.

```
        Output bytes  :               2952
        Input  packets:                  6
        Output packets:                 50
   VPI 10
     Flags: Active
     Total down time: 0 sec, Last down: Never
       Traffic statistics:
        Input  bytes  :                  0
        Output bytes  :                  0
        Input  packets:                  0
        Output packets:                  0

 Logical interface at-0/3/1.0 (Index 78) (SNMP ifIndex 77) (Generation 20)
   Flags: Point-To-Point Copy-PLP-To-CLP SNMP-Traps 0x4000
   Encapsulation: ATM-SNAP
   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: 4470, Generation: 38, Route table: 0
     Flags: None
     Addresses, Flags: Is-Preferred Is-Primary
       Destination: 10.0.59.5, Local: 10.0.59.6, Broadcast: Unspecified,
       Generation: 44
   Protocol iso, MTU: 4470, Generation: 39, Route table: 0
     Flags: None
   VCI 0.128
     Flags: Active
     Total down time: 0 sec, Last down: Never
     EPD threshold: 2129, Transmit weight cells: 10
     ATM per-VC transmit statistics:
      Tail queue packet drops: 0
     Traffic statistics:
      Input  bytes  :                  0
      Output bytes  :                  0
      Input  packets:                  0
      Output packets:                  0
 Logical interface at-0/3/1.32767 (Index 79) (SNMP ifIndex 76) (Generation 21)
   Flags: Point-To-Multipoint Copy-PLP-To-CLP No-Multicast SNMP-Traps 0x4000
   Encapsulation: ATM-VCMUX
   Traffic statistics:
    Input  bytes  :                360
    Output bytes  :               3302
    Input  packets:                  6
    Output packets:                 50
   Local statistics:
    Input  bytes  :                360
    Output bytes  :               3302
    Input  packets:                  6
    Output packets:                 50
```

 Copyright © 2018, Juniper Networks, Inc.

```
    VCI 0.16
      Flags: Active, ILMI
      Total down time: 0 sec, Last down: Never
      EPD threshold: 0, Transmit weight cells: 0
      ATM per-VC transmit statistics:
       Tail queue packet drops: 0
      Traffic statistics:
       Input  bytes  :                    0
       Output bytes  :                 2640
       Input  packets:                    0
       Output packets:                   44
    VCI 0.4
      Flags: Active, OAM
      OAM, Period 30 sec, Up count: 10, Down count: 10
      Total down time: 0 sec, Last down: Never
      EPD threshold: 2129, Transmit weight cells: 0
      ATM per-VC transmit statistics:
       Tail queue packet drops: 0
      Traffic statistics:
       Input  bytes  :                  312
       Output bytes  :                  312
       Input  packets:                    6
       Output packets:                    6
     OAM F4 cell statistics:
      Total received: 6, Total sent: 6
      Loopback received: 6, Loopback sent: 6
      Last received: 00:00:29, Last sent: 00:00:29
      RDI received: 0, RDI sent: 0
      AIS received: 0, AIS sent: 0
```

show interfaces extensive (ATM2, SONET Mode)

```
user@host> show interfaces at-0/3/1 extensive
Physical interface: at-0/3/1, Enabled, Physical link is Up
  Interface index: 139, SNMP ifIndex: 67, Generation: 22
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SONET mode,
  Speed: OC3, Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:5e
  Last flapped   : 2006-03-13 17:46:36 PST (16:04:12 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                  520                    0 bps
   Output bytes  :                 4240                    0 bps
   Input  packets:                   10                    0 pps
   Output packets:                   72                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Invalid VCs: 0, Framing errors: 0, Policed discards: 0,

    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    Resource errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, Aged packets: 0, MTU errors: 0,

    Resource errors: 0
  Egress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
```

Copyright © 2018, Juniper Networks, Inc.

```
  0 best-effort                  62                  62                   0

  1 expedited-fo                  0                   0                   0

  2 assured-forw                  0                   0                   0

  3 network-cont                 10                  10                   0

SONET alarms   : None
SONET defects  : None
SONET PHY:           Seconds        Count  State
  PLL Lock                 0            0  OK
  PHY Light                0            0  OK
SONET section:
  BIP-B1                   0            0
  SEF                      0            0  OK
  LOS                      0            0  OK
  LOF                      0            0  OK
  ES-S                     0
  SES-S                    0
  SEFS-S                   0
SONET line:
  BIP-B2                   0            0
  REI-L                    0            0
  RDI-L                    0            0  OK
  AIS-L                    0            0  OK
  BERR-SF                  0            0  OK
  BERR-SD                  0            0  OK
  ES-L                     0
  SES-L                    0
  UAS-L                    0
  ES-LFE                   0
  SES-LFE                  0
  UAS-LFE                  0
SONET path:
  BIP-B3                   0            0
  REI-P                    0            0
  LOP-P                    0            0  OK
  AIS-P                    0            0  OK
  RDI-P                    0            0  OK
  UNEQ-P                   1            1  OK
  PLM-P                    0            0  OK
  ES-P                     1
  SES-P                    1
  UAS-P                    0
  ES-PFE                   0
  SES-PFE                  0
  UAS-PFE                  0
Received SONET overhead:
  F1      : 0x00, J0      : 0x00, K1      : 0x00, K2      : 0x00
  S1      : 0x00, C2      : 0x13, C2(cmp) : 0x13, F2      : 0x00
  Z3      : 0x00, Z4      : 0x00, S1(cmp) : 0x00
Transmitted SONET overhead:
  F1      : 0x00, J0      : 0x01, K1      : 0x00, K2      : 0x00
  S1      : 0x00, C2      : 0x13, F2      : 0x00, Z3      : 0x00
  Z4      : 0x00
ATM status:
  HCS state:     Sync
  LOC      :       OK
ATM Statistics:
  Uncorrectable HCS errors: 0, Correctable HCS errors: 0,
```


Copyright © 2018, Juniper Networks, Inc.

```
    Tx cell FIFO overruns: 0, Rx cell FIFO overruns: 0,
    Rx cell FIFO underruns: 0, Input cell count: 0, Output cell count: 0,
    Output idle cell count: 0, Output VC queue drops: 0, Input no buffers: 0,
    Input length errors: 0, Input timeouts: 0, Input invalid VCs: 0,
    Input bad CRCs: 0, Input OAM cell no buffers: 0
  Packet Forwarding Engine configuration:
    Destination slot: 0
    VPI 0
      Flags: Active, OAM, Shaping
      CBR, Peak: 50kbps
      OAM, Period 30 sec, Up count: 10, Down count: 10
      Total down time: 0 sec, Last down: Never
     OAM F4 cell statistics:
      Total received: 10, Total sent: 10
      Loopback received: 10, Loopback sent: 10
      Last received: 00:00:02, Last sent: 00:00:02
      RDI received: 0, RDI sent: 0
      AIS received: 0
        Traffic statistics:
         Input  bytes  :                520
         Output bytes  :               4240
         Input  packets:                 10
         Output packets:                 72
    VPI 10
      Flags: Active
      Total down time: 0 sec, Last down: Never
        Traffic statistics:
         Input  bytes  :                  0
         Output bytes  :                  0
         Input  packets:                  0
         Output packets:                  0

  Logical interface at-0/3/1.0 (Index 78) (SNMP ifIndex 77) (Generation 20)
    Flags: Point-To-Point Copy-PLP-To-CLP SNMP-Traps 0x4000
    Encapsulation: ATM-SNAP
    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: 4470, Generation: 38, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.0.59.5, Local: 10.0.59.6, Broadcast: Unspecified,
        Generation: 44
    Protocol iso, MTU: 4470, Generation: 39, Route table: 0
      Flags: None
    VCI 0.128
      Flags: Active
      Total down time: 0 sec, Last down: Never
      EPD threshold: 2129, Transmit weight cells: 10
```

Copyright © 2018, Juniper Networks, Inc.

```
      ATM per-VC transmit statistics:
       Tail queue packet drops: 0
      Traffic statistics:
       Input  bytes  :                   0
       Output bytes  :                   0
       Input  packets:                   0
       Output packets:                   0

  Logical interface at-0/3/1.32767 (Index 79) (SNMP ifIndex 76) (Generation 21)
    Flags: Point-To-Multipoint Copy-PLP-To-CLP No-Multicast SNMP-Traps 0x4000
    Encapsulation: ATM-VCMUX
    Traffic statistics:
     Input  bytes  :                660
     Output bytes  :               5473
     Input  packets:                 11
     Output packets:                 83
    Local statistics:
     Input  bytes  :                660
     Output bytes  :               5473
     Input  packets:                 11
     Output packets:                 83
    VCI 0.16
      Flags: Active, ILMI
      Total down time: 0 sec, Last down: Never
      EPD threshold: 0, Transmit weight cells: 0
      ATM per-VC transmit statistics:
       Tail queue packet drops: 0
      Traffic statistics:
       Input  bytes  :                   0
       Output bytes  :                4320
       Input  packets:                   0
       Output packets:                  72
    VCI 0.4
      Flags: Active, OAM
      OAM, Period 30 sec, Up count: 10, Down count: 10
      Total down time: 0 sec, Last down: Never
      EPD threshold: 2129, Transmit weight cells: 0
      ATM per-VC transmit statistics:
       Tail queue packet drops: 0
      Traffic statistics:
       Input  bytes  :                 572
       Output bytes  :                 572
       Input  packets:                  11
       Output packets:                  11
     OAM F4 cell statistics:
      Total received: 11, Total sent: 11
      Loopback received: 11, Loopback sent: 11
      Last received: 00:00:18, Last sent: 00:00:18
      RDI received: 0, RDI sent: 0
      AIS received: 0, AIS sent: 0
```

 Copyright © 2018, Juniper Networks, Inc.

show interfaces (Channelized DS3-to-DS0)

Syntax

show interfaces ds-*fpc/pic/port:t1channel:ds0channel*
<brief | detail | extensive>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced before Junos OS Release 7.4.

Description

(M Series and T Series routers only) Display status information about the specified channelized DS3-to-DS0 interface.

Options

ds-*fpc/pic/port:t1channel:ds0channel*—Display standard information about the specified channelized DS3-to-DS0 interface.

brief | detail | extensive—(Optional) Display the specified level of output interface.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

List of Sample Output

show interfaces extensive (Channelized DS3-to-DS0) on page 1241

Output Fields

Table 60 on page 1233 lists the output fields for the **show interfaces** (all Channelized DS3 interfaces) command. Output fields are listed in the approximate order in which they appear.

Table 60: Channelized DS3 show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under “Common Output Fields Description” on page 1110.	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

Table 60: Channelized DS3 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**
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**.	All levels
Speed	Speed at which the interface is running.	All levels
Loopback	Whether loopback is enabled and the type of loopback (**local** or **remote**).	All levels
FCS	Frame check sequence on the interface (either **16** or **32**). The default is **16** bits.	All levels
Mode	Whether C-bit parity mode or M13 mode is enabled.	All levels
Framing	Physical layer framing format used on the link. It can be **ESF** or **SF**. The default is **ESF**.	All levels
Parent	(Channelized IQ interfaces only) Name and interface index of the interface to which a particular child interface belongs. **None** indicates that this interface is the top level.	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 1110.	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 1110.	All levels
Link flags	Information about the link. Possible values are described in the “Link Flags” section under “Common Output Fields Description” on page 1110.	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 keepalives. • **interval** *seconds*—The time in seconds between successive keepalive requests. The range is **10** seconds through **32,767** seconds, with a default of **10** seconds. • **down-count** *number*—The number of keepalive packets that a destination must fail to receive before the network takes a link down. The range is **1** through **255**, with a default of **3**. • **up-count** *number*—The number of keepalive packets a destination must receive to change a link’s status from down to up. The range is **1** through **255**, with a default of **1**.	**detail extensive** none

 Copyright © 2018, Juniper Networks, Inc.

Table 60: Channelized DS3 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output

Keepalive statistics	(PPP and HDLC) Information about keepalive packets.
	• **Input**—Number of keepalive packets received by PPP.
	  ◦ **(last seen 00:00:00 ago)**—Time since the last keepalive packet was received, in the format ***hh:mm:ss***.
	• **Output**—Number of keepalive packets sent by PPP and how long ago the last keepalive packets were sent and received.
	  ◦ **(last seen 00:00:00 ago)**—Time since the last keepalive packet was sent, in the format ***hh:mm:ss***.
		detail extensive none

LMI settings	(Frame Relay) Settings for Local Management Interface (LMI) can be either **ANSI LMI settings** or **ITU LMI settings**. ANSI LMI settings is the default. The format is **(ANSI or ITU) LMI settings:** ***value***, *value*, *value...xx* seconds, where ***value*** can be: • **n391dte**—DTE full status polling interval (1–255) • **n392dce**—DCE error threshold (1–10) • **n392dte**—DTE error threshold (1–10) • **n393dce**—DCE monitored event count (1–10) • **n393dte**—DTE monitored event count (1–10) • **t391dte**—DTE polling timer (5–30 seconds) • **t392dce**—DCE polling verification timer (5–30 seconds)	**detail extensive** none
LMI	(Frame Relay) LMI packet statistics: • **Input**—Number of packets coming in on the interface (*nn*) and how much time has passed since the last packet arrived. The format is **Input:** ***nn*** **(last seen** ***hh:mm:ss*** **ago)**. • **Output**—Number of packets sent out on the interface (*nn*) and how much time has passed since the last packet was sent. The format is **Output:** ***nn*** **(last seen** ***hh:mm:ss*** **ago)**.	**detail extensive** none
LCP state	(PPP) Link Control Protocol state. • **Conf-ack-received**—Acknowledgement was received. • **Conf-ack-sent**—Acknowledgement was sent. • **Conf-req-sent**—Request was sent. • **Down**—LCP negotiation is incomplete (not yet completed or has failed). • **Not-configured**—LCP is not configured on the interface. • **Opened**—LCP negotiation is successful.	**detail extensive** none
NCP state	(PPP) Network Control Protocol state. • **Conf-ack-received**—Acknowledgement was received. • **Conf-ack-sent**—Acknowledgement was sent. • **Conf-req-sent**—Request was sent. • **Down**—NCP negotiation is incomplete (not yet completed or has failed). • **Not-configured**—NCP is not configured on the interface. • **Opened**—NCP negotiation is successful.	**detail extensive** none

Table 60: Channelized DS3 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
CHAP state	(PPP) Displays the state of the Challenge Handshake Authentication Protocol (CHAP) during its transaction. • **Chap-Chal-received**—Challenge was received but response not yet sent. • **Chap-Chal-sent**—Challenge was sent. • **Chap-Resp-received**—Response was received for the challenge sent, but CHAP has not yet moved into the**Success** state. (Most likely with RADIUS authentication). • **Chap-Resp-sent**—Response was sent for the challenge received. • **Closed**—CHAP authentication is incomplete. • **Failure**—CHAP authentication failed. • **Not-configured**—CHAP is not configured on the interface. • **Success**—CHAP authentication was successful.	**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 hh:mm:ss*** **ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.	**detail extensive** none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface. • **Input packets, Output packets**—Number of packets received and transmitted on the interface.	**detail extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 60: Channelized DS3 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Giants**—Number of frames received that are larger than the giant threshold. • **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. • **L2 channel errors**—Counter increments when the software could not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Count of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **SRAM errors**—Number of hardware errors that occurred in the static RAM (SRAM) on the PIC. If the value in this field increments, the PIC is malfunctioning. • **HS link CRC errors**—Count of errors on the high-speed links between the ASICs responsible for handling the router interfaces.	**extensive**
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly, (perhaps once every 10 seconds), the cable, the far-end system, or the 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. • **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.	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 60: Channelized DS3 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
DS1 alarms **DS1 defects**	Media-specific defects that can render the interface unable to pass 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. • **LOS**—Loss of signal. • **LOF**—Loss of frame. • **AIS**—Alarm indication signal. • **YLW**—Yellow alarm. Indicates errors at the remote site receiver.	**detail extensive** none
T1 media	Counts of T1 media-specific errors. • **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. The T1 media-specific error types can be: • **SEF**—Severely errored framing • **BEE**—Bit error event • **AIS**—Alarm indication signal • **LOF**—Loss of frame • **LOS**—Loss of signal • **YELLOW**—Errors at the remote site receiver • **BPV**—Bipolar violation • **EXZ**—Excessive zeros • **LCV**—Line code violation • **PCV**—Pulse code violation • **CS**—Carrier state • **LES**—Line error seconds • **ES**—Errored seconds • **SEFS**—Severely errored framing seconds (section) • **SES**—Severely errored seconds • **BES**—Bit error seconds • **UAS**—Unavailable seconds	**extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 60: Channelized DS3 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
DS3 media	Counts of T3 media-specific errors. For detailed definitions of the T3 (DS-3) error events (BPV, EXZ, LCV, PCV, and CCV) and performance parameters (LES, PES, PSES, CES, CSES, SEFS, and UAS), see RFC 2496. • **PLL Lock**—Phase-locked loop out of lock • **Reframing**—Frame alignment recovery time • **AIS**—Alarm indication signal • **LOF**—Loss of frame • **LOS**—Loss of signal • **IDLE**—Idle code detected • **YELLOW**—Remote defect indication • **BPV**—Bipolar violation • **EXZ**—Excessive zeros • **LCV**—Line code violation • **PCV**—Pulse code violation • **CCV**—C-bit coding violation • **LES**—Line error seconds • **PES**—P-bit errored seconds • **PSES**—P-bit errored seconds (section) • **CES**—C-bit errored seconds • **CSES**—C-bit severely errored seconds • **SEFS**—Severely errored framing seconds • **UAS**—Unavailable seconds	**extensive**
HDLC configuration	Information about the HDLC configuration. • **Giant threshold**—Giant threshold programmed into the hardware. • **Runt threshold**—Runt threshold programmed into the hardware. • **Timeslots**—Configured time slots for the interface. • **Byte encoding**—Byte encoding used: **Nx64K** or **Nx56K**. • **Data inversion**—HDLC data inversion setting: **Enabled** or **Disabled**	**extensive**
Interface transmit queues	Name of the transmit queues and their associated statistics for each DS1 channel on the Channelized DS3-to-DS1 PIC. • **B/W**—Queue bandwidth as a percentage of the total interface bandwidth. • **WRR**—Weighted round-robin (in percent). • **Packets**—Number of packets transmitted. • **Bytes**—Number of bytes transmitted. • **Drops**—Number of packets dropped. • **Errors**—Number of packet errors.	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 60: Channelized DS3 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
DS1 or DS3 BERT configuration	BERT (bit error rate test) checks the quality of the line. This output appears only when a BERT is run on the interface. • **BERT time period**—Configured total time period that the BERT is to run. • **Elapsed**—Actual time elapsed since the start of the BERT (in seconds). • **Induced error rate**—Configured rate at which the bit errors are induced in the BERT pattern. • **Algorithm**—Type of algorithm selected for the BERT.	**detail extensive** none
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number. • **PLP byte**—Packet Level Protocol byte.	**extensive**
CoS information	Information about the CoS queue for the physical interface. • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	Logical interface SNMP interface index number.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface; values are described in the "Logical Interface Flags" section under "Common Output Fields Description" on page 1110.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Bandwidth	Bandwidth configured on the interface.	All levels
Protocol	Protocol family configured on the logical interface, such as **iso**, **inet6**, **mpls**.	**detail extensive** none

Copyright © 2018, Juniper Networks, Inc.

Table 60: Channelized DS3 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
MTU	MTU size on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route table	Routing table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	detail extensive
Flags	Information about the protocol family flags. Possible values are described in the “Family Flags” section under “Common Output Fields Description” on page 1110.	**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 1110.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive** none
Redundant Link	(LSQ redundancy) Backup link for Link Services IQ redundancy.	**detail extensive** none

Sample Output

show interfaces extensive (Channelized DS3-to-DS0)

```
user@host> show interfaces ds-0/0/0:0:0 extensive
Physical interface: ds-0/0/0:0:0, Enabled, Physical link is Up
  Interface index: 174, SNMP ifIndex: 4298, Generation: 177
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: 640kbps, FCS: 16,
  Mode: C/Bit parity, Framing: ESF
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  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 : 280 (last seen 00:00:09 ago)
    Output: 286 (last sent 00:00:00 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured,
  mpls: Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-23 17:53:29 PDT (00:46:46 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                 6814                   16 bps
   Output bytes  :                28840                   72 bps
```

Copyright © 2018, Juniper Networks, Inc.

```
 Input  packets:                 568                   0 pps
 Output packets:                 893                   0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 39, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 2, L2 mismatch timeouts: 0,
  HS link CRC errors: 0
Output errors:
  Carrier transitions: 3, Errors: 0, Drops: 0, Aged packets: 0
DS1   alarms   : None
DS3   alarms   : None
DS1   defects  : None
DS3   defects  : None
T1  media:            Seconds        Count  State
  SEF                       0            0  OK
  BEE                       5            1  OK
  AIS                       0            0  OK
  LOF                       0            0  OK
  LOS                       0            0  OK
  YELLOW                   17            1  OK
  BPV                       0            0
  EXZ                       0            0
  LCV                       5        27765
  PCV                       0            0
  CS                        0            0
  LES                       0
  ES                        0
  SES                       5
  SEFS                     10
  BES                       0
  UAS                       0
DS3 media:            Seconds        Count  State
  PLL Lock                  0            0  OK
  Reframing                 0            0  OK
  AIS                       0            0  OK
  LOF                       0            0  OK
  LOS                       0            0  OK
  IDLE                      0            0  OK
  YELLOW                    0            0  OK
  BPV                       1        65535
  EXZ                       1        65535
  LCV                       2       131070
  PCV                       1         1825
  CCV                       0            0
  LES                       1
  PES                       1
  PSES                      1
  CES                       0
  CSES                      0
  SEFS                      0
  UAS                       0
Interface transmit queues:
            B/W  WRR       Packets        Bytes        Drops       Errors
  Queue0     95   95             0            0            0            0
  Queue1      5    5           893        28840            0            0
HDLC configuration:
  Giant threshold: 1514, Runt threshold: 3
  Timeslots      : 1-10
  Byte encoding: Nx64K, Data inversion: Disabled
DS3 BERT configuration:
  BERT time period: 10 seconds, Elapsed: 0 seconds
  Algorithm: 2^15 - 1, Induced error rate: 10e-0
```

 Copyright © 2018, Juniper Networks, Inc.

```
  DS1 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Induced Error rate: 10e-0, Algorithm: 2^15 - 1, 0.151, Pseudorandom (9)
  Packet Forwarding Engine configuration:
    Destination slot: 0, PLP byte: 2 (0x01)
  CoS information:    CoS transmit queue            Bandwidth           Buffer
Priority   Limit
                            %          bps   %         usec
    0 best-effort          95       608000  95            0       low    none
    3 network-control       5        32000   5            0       low    none
 Logical interface ds-0/0/0:0:0.0 (Index 5) (SNMP ifIndex 4299)
  (Generation 943)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 1500, Generation: 949, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 198.51.100.0/24, Local: 198.51.100.1, Broadcast:
198.51.100.255,
        Generation: 1849
```

Copyright © 2018, Juniper Networks, Inc.

show interfaces (Channelized DS3-to-DS1)

Syntax

show interfaces t1-*fpc/pic/port:t1channel*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced before Junos OS Release 7.4.

Description

(M Series and T Series routers only) Display status information about the specified channelized DS3-to-DS1 interface.

Options

t1-*fpc/pic/port:t1channel*—Display standard information about the specified channelized DS3-to-DS1 interface.

brief | detail | extensive | terse—(Optional) Display brief, detail, extensive, or terse information about the interface.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

List of Sample Output

show interfaces extensive (channelized DS3-to-DS1) on page 1244

Output Fields

See the output field table for the show interfaces (Channelized DS3-to-DS0) command.

Sample Output

show interfaces extensive (channelized DS3-to-DS1)

```
user@host> show interfaces t1-0/0/0:0 extensive
Physical interface: t1-0/0/0:0, Enabled, Physical link is Up
  Interface index: 210, SNMP ifIndex: 14, Generation: 2977
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: 640kbps,
  Loopback: None, FCS: 16, Mode: C/Bit parity, Framing: ESF
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  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 : 30 (last seen 00:00:05 ago)
```


Copyright © 2018, Juniper Networks, Inc.

```
  Output: 29 (last sent 00:00:00 ago)
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
CHAP state: Not-configured
Last flapped   : 2002-05-23 17:30:12 PDT (17:29:43 ago)
Statistics last cleared: Never
Traffic statistics:
 Input  bytes  :                944                  16 bps
 Output bytes  :               1162                  16 bps
 Input  packets:                 66                   0 pps
 Output packets:                 82                   0 pps
Input errors:
  Errors: 1, Drops: 0, Framing errors: 1, Policed discards: 8,
  L3 incompletes: 0, L2 channel errors: 1, L2 mismatch timeouts: 0,
  HS link CRC errors: 0, SRAM errors: 0
Output errors:
  Carrier transitions: 3, Errors: 0, Drops: 0, Aged packets: 0
DS1   alarms   : None
DS3   alarms   : None
DS1   defects  : None
DS3   defects  : None
T1  media:            Seconds        Count  State
  SEF                       0            0  OK
  BEE                      11            5  OK
  AIS                      28            1  OK
  LOF                      27            1  OK
  LOS                       0            0  OK
  YELLOW                   23            1  OK
  BPV                       0            0
  EXZ                       0            0
  LCV                      11        20574
  PCV                       0            0
  CS                        0            0
  LES                      28
  ES                       28
  SES                      39
  SEFS                     50
  BES                       0
  UAS                       0
DS3 media:            Seconds        Count  State
  PLL Lock                  0            0  OK
  Reframing                 0            0  OK
  AIS                       0            0  OK
  LOF                       1            1  OK
  LOS                       1            1  OK
  IDLE                      0            0  OK
  YELLOW                    0            0  OK
  BPV                       2       131070
  EXZ                       3        49910
  LCV                       5       180980
  PCV                       2          327
  CCV                      12       264558
  LES                       3
  PES                       3
  PSES                      2
  CES                      13
  CSES                     13
  SEFS                      1
  UAS                      35
Interface transmit queues:
```

```
             B/W  WRR       Packets        Bytes        Drops       Errors
    Queue0     95   95             0            0            0            0
    Queue1      5    5            82         1162            0            0
  HDLC configuration:
    Giant threshold: 1514, Runt threshold: 3
    Timeslots      : 1-10
    Line encoding: B8ZS, Byte encoding: Nx64K, Data inversion: Disabled
  DS3 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Algorithm: 2^15 - 1, Induced error rate: 10e-0
  DS1 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Induced Error rate: 10e-0, Algorithm: 2^15 - 1, 0.151, Pseudorandom (9)
  Packet Forwarding Engine configuration:
    Destination slot: 0, PLP byte: 2 (0x00)  CoS information:
    CoS transmit queue             Bandwidth           Buffer Priority   Limit
                              %          bps   %         usec
    0 best-effort            95       608000  95            0      low    none
    3 network-control         5        32000   5            0      low    none
  Logical interface t1-0/0/0:0.0 (Index 11) (SNMP ifIndex 23) (Generation 497)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Bandwidth: 0
    Protocol inet, MTU: 1500, Generation: 576, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 198.51.100.0/24, Local: 198.51.100.1, Broadcast:
198.51.100.255,
        Generation: 977
```


Copyright © 2018, Juniper Networks, Inc.

show interfaces (Channelized E1 IQ)

Syntax

show interfaces (ce1-*fpc/pic/port* | *type-fpc/pic/port*<:*channel*>)
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced before Junos OS Release 7.4.

Description

(M Series and T Series routers only) Display status information about the specified channelized E1 IQ interface.

Options

type-fpc/pic/port*:<*channel*>**—Interface type with optional corresponding channel levels. For the physical channelized E1 IQ interface, ***type is **ce**. For the clear channel, ***type*** is **e1**. At the first level of channelization, ***type*** is **ds**.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

List of Sample Output

show interfaces (Channelized E1 IQ) (Physical) on page 1247
show interfaces extensive (Channelized E1 IQ Multilink PPP Encapsulation) on page 1248
show interfaces extensive (Channelized E1 IQ MLFR Encapsulation) on page 1249
show interfaces detail (Clear Channel E1) on page 1250

Output Fields

For information about output fields, see the output field table for the show interfaces (Channelized E1) command. Output fields are listed in the approximate order in which they appear.

Sample Output

show interfaces (Channelized E1 IQ) (Physical)

```
user@host> show interfaces ce1-1/2/3
Physical interface: ce1-1/2/3, Enabled, Physical link is Up
  Interface index: 18, SNMP ifIndex: 1128
  Link-level type: Frame-relay, Controller, MTU: 1504, Clocking: Internal,   Speed:
 E1, Loopback: None, FCS: 16, Framing: G704, Parent: None
```

Copyright © 2018, Juniper Networks, Inc.

```
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives DTE
  ANSI LMI settings: n391dte 6, n392dte 3, n393dte 4, t391dte 10 seconds
  LMI: Input: 51700 (00:00:02 ago), Output: 51701 (00:00:02 ago)
  DTE statistics:
    Enquiries sent                     : 43186
    Full enquiries sent                : 8515
    Enquiry responses received         : 43185
    Full enquiry responses received    : 8515
  DCE statistics:
    Enquiries received                 : 0
    Full enquiries received            : 0
    Enquiry responses sent             : 0
    Full enquiry responses sent        : 0
  Common statistics:
    Unknown messages received          : 0
    Asynchronous updates received      : 0
    Out-of-sequence packets received   : 0
    Keepalive responses timedout       : 0
  Nonmatching DCE-end DLCIs:
          2
  Last flapped   : 2002-10-04 17:52:51 PDT (00:32:57 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  DS1   alarms   : None
  DS1   defects  : None
```

show interfaces extensive (Channelized E1 IQ Multilink PPP Encapsulation)

```
user@host> show interfaces ds-0/3/4:1 extensive
Physical interface: ds-0/3/4:1, Enabled, Physical link is Up
  Interface index: 151, SNMP ifIndex: 63, Generation: 34
  Link-level type: Multilink-PPP, MTU: 1518, Clocking: Internal, Speed: 64kbps,
  Loopback: None, FCS: 16,
  Parent: ce1-0/3/4 Interface index 150
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Link flags     : None
  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
  CHAP state: Closed
  CoS queues     : 4 supported, 4 maximum usable queues
  Last flapped   : Never
  Statistics last cleared: 2005-12-21 10:32:15 PST (1w0d 03:10 ago)
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :              6070570                  224 bps
   Input  packets:                    0                    0 pps
   Output packets:               209330                    0 pps
  Input errors:
    Errors: 3, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, L3 incompletes: 0, L2 channel errors: 0,
    L2 mismatch timeouts: 0, HS link CRC errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Aged packets: 0, MTU errors: 0,
     Resource errors: 0
```

 Copyright © 2018, Juniper Networks, Inc.

```
HDLC configuration:
  Giant threshold: 1528, Runt threshold: 2
  Timeslots      : 1
  Data inversion: Disabled, Idle cycle flag: flags, Start end flag: shared
DS0 BERT configuration:
  BERT time period: 10 seconds, Elapsed: 0 seconds
  Induced Error rate: 10e-0, Algorithm: 2^15 - 1, 0.151, Pseudorandom (9)
Packet Forwarding Engine configuration:
  Destination slot: 0, PLP byte: 4 (0x00)

Logical interface ds-0/3/4:1.0 (Index 74) (SNMP ifIndex 64) (Generation 13)
  Flags: Hardware-Down Point-To-Point SNMP-Traps Encapsulation: PPP
  Protocol mlppp, Multilink bundle: ls-0/1/0.0, MTU: 1514, Generation: 24,
  Route table: 0
```

show interfaces extensive (Channelized E1 IQ MLFR Encapsulation)

```
user@host> show interfaces ds-0/3/4:5 extensive
Physical interface: ds-0/3/4:5, Enabled, Physical link is Up
  Interface index: 155, SNMP ifIndex: 72, Generation: 38
  Link-level type: Multilink-FR, MTU: 1518, Clocking: Internal, Speed: 64kbps,
  Loopback: None, FCS: 16,
  Parent: ce1-0/3/4 Interface index 150
  Device flags   : Present Running
  Interface flags: Link-Layer-Down Point-To-Point SNMP-Traps Internal: 0x4000
  Link flags     : No-Keepalives DCE
  Hold-times     : Up 0 ms, Down 0 ms
  ANSI LMI settings: n392dce 3, n393dce 4, t392dce 15 seconds
  LMI statistics:
    Input : 0 (last seen: never)
    Output: 0 (last sent: never)
  DTE statistics:
    Enquiries sent                      : 0
    Full enquiries sent                 : 0
    Enquiry responses received          : 0
    Full enquiry responses received     : 0
  DCE statistics:
    Enquiries received                  : 0
    Full enquiries received             : 0
    Enquiry responses sent              : 0
    Full enquiry responses sent         : 0
  Common statistics:
    Unknown messages received           : 0
    Asynchronous updates received       : 0
    Out-of-sequence packets received    : 0
    Keepalive responses timedout        : 0
  CoS queues     : 4 supported, 4 maximum usable queues
  Last flapped   : 2005-12-21 09:59:01 PST (1w0d 03:44 ago)
  Statistics last cleared: 2005-12-21 10:32:15 PST (1w0d 03:10 ago)
  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: 3, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, L3 incompletes: 0, L2 channel errors: 0,
    L2 mismatch timeouts: 0, HS link CRC errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Aged packets: 0, MTU errors: 0,
     Resource errors: 0
```

Copyright © 2018, Juniper Networks, Inc.

```
  HDLC configuration:
    Giant threshold: 1528, Runt threshold: 2
    Timeslots      : 5
    Data inversion: Disabled, Idle cycle flag: flags, Start end flag: shared
  DS0 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Induced Error rate: 10e-0, Algorithm: 2^15 - 1, 0.151, Pseudorandom (9)
  Packet Forwarding Engine configuration:
    Destination slot: 0, PLP byte: 4 (0x01)

  Logical interface ds-0/3/4:5.0 (Index 78) (SNMP ifIndex 73) (Generation 17)
    Flags: Device-Down Point-To-Point SNMP-Traps Encapsulation: FR-NLPID
    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 mlfr, Multilink bundle: ls-0/1/0.1, MTU: 1514, Generation: 28, Route
 table: 0
    DLCI 10
      Flags: Active
      Total down time: 0 sec, Last down: Never
      Traffic statistics:
       Input  bytes  :                    0
       Output bytes  :                    0
       Input  packets:                    0
       Output packets:                    0
  DLCI statistics:
    Active DLCI   :1  Inactive DLCI   :0
```

show interfaces detail (Clear Channel E1)

```
user@host> show interfaces e1-1/2/6 detail
Physical interface: e1-1/2/6, Enabled, Physical link is Up
  Interface index: 89, SNMP ifIndex: 1278, Generation: 341
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: E1, Loopback:None,
...
  Logical interface e1-1/2/6.0 (Index 52) (SNMP ifIndex 1279) (Generation 169)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Bandwidth: 0
...
```

 Copyright © 2018, Juniper Networks, Inc.

show interfaces (Channelized E1)

Syntax

show interfaces ds-*fpc/pic/port:ds0channel*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information Command introduced before Junos OS Release 7.4.

Description (M Series and T Series routers only) Display status information the specified channelized E1 interface.

Options

ds-*fpc/pic/port:ds0channel*—Display standard information about the specified channelized E1 interface.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level view

List of Sample Output show interfaces extensive (Channelized E1) on page 1260

Output Fields Table 61 on page 1251 lists the output fields for the **show interfaces** (Channelized E1 and Channelized E1 IQ) command. Output fields are listed in the approximate order in which they appear.

Table 61: Channelized E1 and Channelized E1 IQ show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under “Common Output Fields Description” on page 1110.	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

Table 61: Channelized E1 and Channelized E1 IQ 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**
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: **Internal** or **External**.	All levels
Speed	Speed at which the interface is running.	All levels
Loopback	Whether loopback is enabled and the type of loopback (**local** or **remote**).	All levels
FCS	Frame check sequence on the interface (either **16** or **32**). The default is **16** bits.	All levels
Framing	Physical layer framing format used on the link. It can be **G704**, **G704-NO-CRC4**, or **Unframed**. The default is **G704**.	All levels
Parent	(Channelized E1 IQ interfaces only) Name and interface index of the interface to which a particular child interface belongs. **None** indicates that this interface is the top level.	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 1110.	All levels
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under "Common Output Fields Description" on page 1110.	All levels
Link flags	Information about the link. Possible values are described in the "Link Flags" section under "Common Output Fields Description" on page 1110.	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 keepalives. • **Interval** ***seconds***—Time in seconds between successive keepalive requests. The range is **10** seconds through **32,767** seconds, with a default of **10** seconds. • **Down-count** ***number***—Number of keepalive packets a destination must fail to receive before the network takes a link down. The range is **1** through **255**, with a default of **3**. • **Up-count** ***number***—Number of keepalive packets a destination must receive to change a link's status from down to up. The range is **1** through **255**, with a default of **1**.	**detail extensive** none

 Copyright © 2018, Juniper Networks, Inc.

Table 61: Channelized E1 and Channelized E1 IQ show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output

Keepalive statistics	(PPP and HDLC) Information about keepalive packets.
	• **Input**—Number of keepalive packets received by PPP.
	  ◦ (**last seen 00:00:00 ago**)—Time since the last keepalive packet was received, in the format ***hh:mm:ss***.
	• **Output**—Number of keepalive packets sent by PPP and how long ago the last keepalive packets were sent and received.
	  ◦ (**last seen 00:00:00 ago**)—Time since the last keepalive packet was sent, in the format ***hh:mm:ss***.
		detail extensive none

LMI settings	(Frame Relay) Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is **(ANSI or ITU) LMI settings: *value*, *value*... *xx*** seconds, where ***value*** can be: • **n391dte**—DTE full status polling interval (1–255) • **n392dce**—DCE error threshold (1–10) • **n392dte**—DTE error threshold (1–10) • **n393dce**—DCE monitored event count (1–10) • **n393dte**—DTE monitored event count (1–10) • **t391dte**—DTE polling timer (5–30 seconds) • **t392dce**—DCE polling verification timer (5–30 seconds)	**detail extensive** none
LMI	(Frame Relay) Statistics about the link management. • **Input**—Number of packets coming in on the interface (*nn*) and how much time has passed since the last packet arrived. The format is **Input: *nn* (last seen *hh:mm:ss* ago)** • **Output**—Number of packets sent out on the interface (*nn*) and how much time has passed since the last packet was sent. The format is **Output: *nn* (last sent *hh:mm:ss* ago)**.	**detail extensive** none
DTE statistics	(Frame Relay) Statistics about messages transmitted from the data terminal equipment (DTE) to the data circuit-terminating equipment (DCE): • **Enquiries sent**—Number of link status enquiries sent from the DTE to the DCE. • **Full enquiries sent**—Number of full enquiries sent from the DTE to the DCE. • **Enquiry responses received**—Number of enquiry responses received by the DTE from the DCE. • **Full enquiry responses received**—Number of full enquiry responses sent from the DTE to the DCE.	**detail extensive** none
DCE statistics	(Frame Relay) Statistics about messages transmitted from the DCE to the DTE: • **Enquiries received**—Number of enquiries received by the DCE from the DTE. • **Full enquiries received**—Number of full enquiries received by the DCE from the DTE. • **Enquiry responses sent**—Number of enquiry responses sent from the DCE to the DTE. • **Full enquiry responses sent**—Number of full enquiry responses sent from the DCE to the DTE.	**detail extensive** none

Copyright © 2018, Juniper Networks, Inc.

Table 61: Channelized E1 and Channelized E1 IQ show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Common statistics	(Frame Relay) Statistics about messages sent between the DTE and the DCE: • **Unknown messages received**—Number of received packets that do not fall into any category. • **Asynchronous updates received**—Number of link status peer changes received. • **Out-of-sequence packets received**—Number of packets for which the sequence of the packets received is different from the expected sequence. • **Keepalive responses timedout**—Number of keepalive responses that timed out when no LMI packet was reported for**n392dte** or **n393dce** intervals. (See **LMI settings**).	**detail extensive** none
Nonmatching DCE-end DLCIs	(Frame Relay, displayed only from the DTE) Number of DLCIs configured from the DCE.	**detail extensive** none
LCP state	(PPP) Link Control Protocol state. • **Conf-ack-received**—Acknowledgement was received. • **Conf-ack-sent**—Acknowledgement was sent. • **Conf-req-sent**—Request was sent. • **Down**—LCP negotiation is incomplete (not yet completed or has failed). • **Not-configured**—LCP is not configured on the interface. • **Opened**—LCP negotiation is successful.	**detail extensive** none
CHAP state	(PPP) Displays the state of the Challenge Handshake Authentication Protocol (CHAP) during its transaction. • **Chap-Chal-received**—Challenge was received but response not yet sent. • **Chap-Chal-sent**—Challenge was sent. • **Chap-Resp-received**—Response was received for the challenge sent, but CHAP has not yet moved into the Success state. (Most likely with RADIUS authentication.) • **Chap-Resp-sent**—Response was sent for the challenge received. • **Closed**—CHAP authentication is incomplete. • **Failure**—CHAP authentication failed. • **Not-configured**—CHAP is not configured on the interface. • **Success**—CHAP authentication was successful.	**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
CoS Queues	Number of CoS queues configured.	**detail extensive** none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 61: Channelized E1 and Channelized E1 IQ show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Traffic statistics	Number of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Giants**—Number of frames received that are larger than the giant threshold. • **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. • **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. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 61: Channelized E1 and Channelized E1 IQ show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), then either 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. • **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. • **MTU errors**—Number of packets larger than the MTU threshold. • **Resource errors**—Sum of transmit drops.	**extensive**
DS1 alarms **DS1 defects**	E1 media-specific defects that can render the interface unable to pass 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. The following lists all possible alarms and defects. For complete explanations of most of these alarms and defects, see *Bellcore Telcordia GR-499-CORE*. • **LOS**—Loss of signal. • **LOF**—Loss of frame. • **AIS**—Alarm indication signal. • **YLW**—Yellow alarm. Indicates errors at the remote site receiver.	**detail extensive** none

 Copyright © 2018, Juniper Networks, Inc.

Table 61: Channelized E1 and Channelized E1 IQ show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
E1 media	Active alarms and defects, plus counts of specific E1 errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than OK indicates a problem. The E1 media-specific error types can be: • **SEF**—Severely errored framing • **BEE**—Bit error • **AIS**—Alarm indication signal • **LOF**—Loss of frame • **LOS**—Loss of signal • **YELLOW**—Errors at the remote site receiver • **BPV**—Bipolar violation • **EXZ**—Excessive zeros • **LCV**—Line code violation • **PCV**—Pulse code violation • **CS**—Carrier state • **FEBE**—Far-end block error • **LES**—Line error seconds • **ES**—Errored seconds • **BES**—Bursty errored seconds • **SES**—Severely errored seconds • **SEFS**—Severely errored framing seconds • **UAS**—Unavailable seconds	**extensive**
HDLC configuration	Information about the HDLC configuration. • **Giant threshold**—Giant threshold programmed into the hardware. • **Runt threshold**—Runt threshold programmed into the hardware. • **Timeslots**—Configured time slots for the interface. • **Line encodingHDB3**—Line encoding used.	**extensive**
Interface transmit queues	Names of the transmit queues and their associated statistics for each DS0 channel on the Channelized E1 to DS0 PIC. • **B/W**—Queue bandwidth as a percentage of the total interface bandwidth. • **WRR**—Weighted round robin (in percent). • **Packets**—Number of packets transmitted. • **Bytes**—Number of bytes transmitted. • **Drops**—Number of packets dropped. • **Errors**—Number of packet errors.	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 61: Channelized E1 and Channelized E1 IQ show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
DS*x* BERT configuration	BERT (bit error rate test) checks the quality of the line. This output appears only when a BERT is run on the interface. • **BERT time period**—Configured total time period that the BERT is to run. • **Elapsed**—Actual time elapsed since the start of the BERT (in seconds). • **Induced error rate**—Configured rate at which the bit errors are induced in the BERT pattern. • **Algorithm**—Type of algorithm selected for the BERT.	**detail extensive** none
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number. • **PLP byte**—Packet Level Protocol byte.	**extensive**
CoS information	Information about the CoS queue for the physical interface. • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	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 1110.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Input packets	Number of packets received on the logical interface.	None specified

Copyright © 2018, Juniper Networks, Inc.

Table 61: Channelized E1 and Channelized E1 IQ show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Output packets	Number of packets transmitted on the logical interface.	None specified
Traffic statistics	Number and rate of bytes and packets received and transmitted on the logical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Local statistics	(Frame Relay) Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than one second) for this counter to stabilize.	**detail extensive**
Transit statistics	(Frame Relay) Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	**detail extensive**
Protocol	Protocol family configured on the logical interface, such as **iso**, **inet6**, **mpls**.	**detail extensive** none
Multilink bundle	(Multilink) Interface name for the multilink bundle, if configured.	**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**
DLCI	(Frame Relay) DLCI number of the logical interface. The following DLCI information is displayed: **Flags**, **Total down time**, **Last down**, and **Traffic statistics**. **Flags** is one or more of the following: • **Active**—Set when the link is active and the DTE and DCE are exchanging information. • **Down**—Set when link is active, but no information is received from the DCE. • **Unconfigured**—Set when the corresponding DLCI in the DCE is not configured. • **Configured**—Set when the corresponding DLCI in the DCE is configured. • **Dce-configured**—Displayed when the command is issued from the DTE.	**detail extensive** none
DLCI statistics	(Frame Relay) Data-link connection identifier (DLCI) statistics. • **Active DLCI**—Number of active DLCIs. • **Inactive DLCI**—Number of inactive DLCIs.	**detail extensive** none

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show interfaces extensive (Channelized E1)

```
user@host> show interfaces ds-0/1/1:1 extensive
Physical interface: ds-0/1/1:1, Enabled, Physical link is Down
  Interface index: 163, SNMP ifIndex: 37, Generation: 46
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: E1,
  Loopback: None, FCS: 16, Framing: G704
  Device flags   : Present Running Down
  Interface flags: Hardware-Down Point-To-Point SNMP-Traps Internal: 0x4000
  Link flags     : Keepalives
  Hold-times     : Up 0 ms, Down 0 ms
  CoS queues     : 4 supported, 4 maximum usable queues
  Last flapped   : 2005-12-28 14:44:06 PST (00:00:30 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, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    HS link CRC errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Aged packets: 0, MTU errors: 0,
     Resource errors: 0
  DS1   alarms   : LOF, LOS
  DS1   defects  : LOF, LOS
  E1  media:            Seconds        Count  State
    SEF                  982318            1  Defect Active
    BEE                       0            0  OK
    AIS                       0            0  OK
    LOF                  982318            1  Defect Active
    LOS                  982318            1  Defect Active
    YELLOW                    0            0  OK
    BPV                       1            1
    EXZ                       1            1
    LCV                       1            1
    PCV                       1            2
    CS                        0            0
    FEBE                      1            9
    LES                       1
    ES                   982318
    SES                  982318
    SEFS                 982318
    BES                       1
    UAS                       0
  Interface transmit queues:
              B/W  WRR       Packets        Bytes        Drops       Errors
    Queue0     95   95             0            0            0            0
    Queue1      5    5             0            0            0            0
  HDLC configuration:
    Giant threshold: 1514, Runt threshold: 3
    Timeslots       : 31
    Line encoding: HDB3, Data inversion: Disabled, Idle cycle flag: flags,
    Start end flag: shared
  DS1 BERT configuration:
    BERT time period: 0 seconds, Elapsed: 0 seconds
```

 Copyright © 2018, Juniper Networks, Inc.

```
  Induced Error rate: 10e-0, Algorithm: 2^11 - 1, 0.152 and 0.153 (2047 type),
   Pseudorandom (8)
Packet Forwarding Engine configuration:
  Destination slot: 0, PLP byte: 2 (0x1b)
CoS information:
  CoS transmit queue     Bandwidth        Buffer   Priority   Limit
                   %          bps     %   usec
  0 best-effort     95      1945600    95      0          low    none
  3 network-control  5       102400     5      0          low    none
```

show interfaces (Channelized OC12 IQ and IQE)

Syntax

show interfaces (*type-fpc/pic/port*<:*channel*><:*channel*><:*channel*>)
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display status information about the specified channelized OC12 IQ and IQE interface.

Options

type-fpc/pic/port*:*channel*:*channel*:*channel—Interface type with optional corresponding channel levels.

For SONET mode, the interface type can be one of the following:

- ***type-fpc/pic/port***—For the physical channelized OC12 IQ or IQE interface, ***type*** is **coc12**. For the clear channel, ***type*** is **so** (for OC12).
- ***type-fpc/pic/port*:*channel***—At the first level of channelization, ***type*** can be **coc1**(channelized OC1), **ct3** (from **coc1)**, **so** (for OC3), or **t3.**
- ***type-fpc/pic/port*:*channel*:*channel***—At the second level of channelization, ***type*** can be **ct1** (from **ct3**or **coc1)** or **t1** (from **ct3** or **coc1)**.
- ***type-fpc/pic/port*:*channel*:*channel*:*channel***—At the third level of channelization, ***type*** is **ds** (from **ct1**).

For SDH mode, the interface type can be one of the following:

- ***type-fpc/pic/port***—For the physical channelized OC12 IQ or IQE interface, ***type*** is **cstm4**. For the clear channel, ***type*** is **so** (for SONET/SDH (vc-4-4c)).
- ***type-fpc/pic/port*:*channel***—At the first level of channelization,***type*** can be **so** (from **cstm4)** or **cau4** (from **cstm4)**.
- ***type-fpc/pic/port*:*channel*:*channel***—At the second level of channelization, ***type*** can be **ct3**or **t3** (from or **cau4)**.
- ***type-fpc/pic/port*:*channel*:*channel*:*channel***—At the third level of channelization, ***type*** is **ct1** or **t1** (from **ct3**).
- ***type-fpc/pic/port*:*channel*:*channel*:*channel***—At the fourth level of channelization, ***type*** is **ds** (from **ct1**).

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

Copyright © 2018, Juniper Networks, Inc.

snmp-index ***snmp-index***—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level view

List of Sample Output

Output Fields See the output field table for the show interfaces (Channelized OC3 IQ and IQE) command.

Sample Output

show interfaces extensive (CAU4 on Channelized OC-12 IQ)

```
user@host> show interfaces cau4-0/2/0:1 extensive
Physical interface: cau4-0/2/0:1, Enabled, Physical link is Up
   Interface index: 219, SNMP ifIndex: 139, Generation: 221
   Link-level type: Controller, Clocking: Internal, SDH mode, Speed: OC3,
   Loopback: None, Parent: cstm4-0/2/0 Interface index 216
   Device flags   : Present Running
   Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
   Link flags     : None
...
```

show interfaces extensive (Channelized OC1 on Channelized OC12 IQ)

```
user@host> show interfaces extensive coc1-4/2/0:7
Physical interface: coc1-4/2/0:7, Enabled, Physical link is Up
  Interface index: 381, SNMP ifIndex: 2524, Generation: 728
  Link-level type: Controller, MTU: 4474, Clocking: Internal, SONET mode,
  Speed: 51840kbps, Loopback: None,
  FCS: 16, Payload scrambler: Disabled, Parent: coc12-4/2/0 (Index 266)
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : None
...
```

show interfaces extensive (Channelized OC12 IQ) (Physical)

```
user@host> show interfaces extensive coc12-4/2/0
Physical interface: coc12-4/2/0, Enabled, Physical link is Up
  Interface index: 266, SNMP ifIndex: 1269, Generation: 601
  Link-level type: Controller, MTU: 4474, Clocking: Internal, SONET mode,
```

```
  Speed: OC12, Loopback: None,
  FCS: 16, Payload scrambler: Disabled, Parent: None  Device flags   : Present
Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives DTE
...
```

show interfaces extensive (Channelized T1 from Channelized OC12 IQ)

```
user@host> show interfaces extensive ct1-4/2/0:7:1
Physical interface: ct1-4/2/0:4:1, Enabled, Physical link is Up
  Interface index: 305, SNMP ifIndex: 2410, Generation: 640
  Link-level type: Controller, MTU: 1504, Clocking: Internal, Speed: T1,
  Loopback: None, FCS: 16,
  Framing: ESF, Parent: coc1-4/2/0:7 (Index 304)
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : None
...
```

show interfaces extensive (Channelized T3 on Channelized OC12 IQ)

```
user@host> show interfaces ct3-0/2/0:1 extensive
Physical interface: ct3-0/2/0:1:1, Enabled, Physical link is Up
   Interface index: 220, SNMP ifIndex: 140, Generation: 222
   Link-level type: Controller, Clocking: Internal, Speed: T3, Loopback: None,
   Mode: C/Bit parity, Parent: cau4-0/2/0:1 Interface index 219
   Device flags   : Present Running
   Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
   Link flags     : None
...
```

show interfaces extensive (CSTM4 on Channelized OC-12 IQ)

```
user@host> show interfaces cstm4-0/2/0 extensive
 Physical interface: cstm4-0/2/0, Enabled, Physical link is Up
   Interface index: 216, SNMP ifIndex: 33, Generation: 218
   Link-level type: Controller, Clocking: Internal, SDH mode, Speed: OC12,
   Loopback: None, Parent: None   Device flags   : Present Running
   Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
   Link flags     : None
...
```

show interfaces extensive (DS0 on Channelized OC12 IQ)

```
user@host> show interfaces extensive ds-4/2/0:7:1:1
Physical interface: ds-4/2/0:4:1:1, Enabled, Physical link is Up
  Interface index: 306, SNMP ifIndex: 2411, Generation: 641
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: 64kbps,
  Loopback: None, FCS: 16, Parent: ct1-4/2/0:7:1 (Index 305)
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
...
```

show interfaces extensive (SONET Interface on Channelized OC12 IQ)

```
user@host> show interfaces so-0/2/0:1 extensive
```

 Copyright © 2018, Juniper Networks, Inc.

```
Physical interface: so-0/2/0:1, Enabled, Physical link is Up
  Interface index: 750, SNMP ifIndex: 23, Generation: 11709
  Link-level type: Multilink-FR, MTU: 4474, Clocking: Internal, SONET mode,
  Speed: OC3, Loopback: None, FCS: 16,
  Payload scrambler: Enabled, Parent: coc12-0/2/0 Interface index 749
  Device flags    : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Link flags      : Keepalives DTE
...
```

show interfaces extensive (T1 on Channelized OC12 IQ)

```
user@host> show interfaces t1-0/2/0:1:1:1 extensive
   Physical interface: t1-0/2/0:1:1:1, Enabled, Physical link is Up
   Interface index: 222, SNMP ifIndex: 143, Generation: 226
   Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: T1,
   Loopback: None, FCS: 16, Framing: ESF, Parent: ct3-0/2/0:1:1
   Interface index 221
   Device flags    : Present Running
   Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
   Link flags      : Keepalives
...
```

show interfaces (Channelized OC12)

Syntax

show interfaces t3-*fpc/pic/port*:*t3channel*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display status information about the specified channelized OC12 interface.

Options

t3-*fpc/pic/port*:*t3channel*—Display standard information about the specified channelized OC12 interface.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

List of Sample Output

show interfaces extensive (Channelized OC12) on page 1266

Output Fields

See the output field table for the show interfaces (Channelized OC3 IQ and IQE) command.

Sample Output

show interfaces extensive (Channelized OC12)

```
user@host> show interfaces t3-0/3/0:0 extensive
Physical interface: t3-0/3/0:0, Enabled, Physical link is Up
  Interface index: 32, SNMP ifIndex: 21, Generation: 2719
  Link-level type: Frame-Relay, PPP, MTU: 4474, Clocking: Internal, SONET mode,
  Speed: T3, Loopback: None, SONET Loopback: None, FCS: 16, Mode: C/Bit parity
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives DTE
  ANSI LMI settings: n391dte 6, n392dte 3, n393dte 4, t391dte 10 seconds
  LMI: Input: 51700 (00:00:02 ago), Output: 51701 (00:00:02 ago)
  DTE statistics:
    Enquiries sent                    : 43186
    Full enquiries sent               : 8515
```


Copyright © 2018, Juniper Networks, Inc.

```
    Enquiry responses received        : 43185
    Full enquiry responses received   : 8515
  DCE statistics:
    Enquiries received                : 0
    Full enquiries received           : 0
    Enquiry responses sent            : 0
    Full enquiry responses sent       : 0
  Common statistics:
    Unknown messages received         : 0
    Asynchronous updates received     : 0
    Out-of-sequence packets received  : 0
    Keepalive responses timedout      : 0
  Nonmatching DCE-end DLCIs:
         2
  Hold-times     : Up 0 ms, Down 0 ms
  Last flapped   : 2002-05-23 16:59:03 PDT (18:23:58 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :               1700                   0 bps
   Output bytes  :               1714                   0 bps
   Input  packets:                123                   0 pps
   Output packets:                124                   0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 1100817, Bucket drops: 0,
    Policed discards: 0, L3 incompletes: 0, L2 channel errors: 0,
    L2 mismatch timeouts: 0, HS link CRC errors: 0, SRAM errors: 0
  Output errors:
    Carrier transitions: 3, Errors: 0, Drops: 0, Aged packets: 0
  DS3   alarms   : None
  SONET alarms   : None
  DS3   defects  : None
  SONET defects  : None
  DS3 media:           Seconds        Count  State
    AIS                      0            0  OK
    LOF                     18            1  OK
    LOS                      0            0  OK
    IDLE                     0            0  OK
    YELLOW                   0            0  OK
    BPV                      0            0
    EXZ                      0            0
    LCV                      0            0
    PCV                     36       122399
    CCV                     72        91948
    LES                      0
    PES                     18
    PSES                    18
    CES                     18
    CSES                    18
    SEFS                    18
    UAS                      0
  HDLC configuration:
    Policing bucket: Disabled
    Shaping bucket : Disabled
    Giant threshold: 4484, Runt threshold: 3
  DSU configuration:
    Compatibility mode: None, Scrambling: Disabled, Subrate: Disabled
    FEAC loopback: Inactive, Response: Disabled, Count: 0
  DS3 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Algorithm: Unknown (0), Induced error rate: 10e-0
Interface transmit queues:
```

```
             B/W  WRR       Packets        Bytes        Drops       Errors
  Queue0     95   95             0            0            0            0
  Queue1      5    5           529         6348            0            0
SONET PHY:             Seconds        Count  State
  PLL Lock                   0            0  OK
  PHY Light                 20            1  OK
SONET section:
  BIP-B1                     0            0
  SEF                       20            1  OK
  LOS                       20            1  OK
  LOF                       20            1  OK
  ES-S                      20
  SES-S                     20
  SEFS-S                    20
SONET line:
  BIP-B2                     0            0
  REI-L                      0            0
  RDI-L                      0            0  OK
  AIS-L                      0            0  OK
  BERR-SF                   18            1  OK
  BERR-SD                    2            1  OK
  ES-L                      20
  SES-L                     20
  UAS-L                     10
  ES-LFE                     0
  SES-LFE                    0
  UAS-LFE                    0
SONET path:
  BIP-B3                     0            0
  REI-P                      0            0
  LOP-P                     20            1  OK
  AIS-P                      0            0  OK
  RDI-P                      0            0  OK
  UNEQ-P                     0            0  OK
  PLM-P                     20            1  OK
  ES-P                      20
  SES-P                     20
  UAS-P                     10
  ES-PFE                     0
  SES-PFE                    0
  UAS-PFE                    0
Received SONET overhead:
  F1      : 0x00, J0      : 0x00, K1      : 0x00, K2      : 0x00
  S1      : 0x00, C2      : 0x04, C2(cmp) : 0x04, F2      : 0x00
  Z3      : 0x00, Z4      : 0x00, S1(cmp) : 0x00, V5      : 0x00
  V5(cmp) : 0x00
Transmitted SONET overhead:
  F1      : 0x00, J0      : 0x01, K1      : 0x00, K2      : 0x00
  S1      : 0x00, C2      : 0x04, F2      : 0x00, Z3      : 0x00
  Z4      : 0x00, V5      : 0x00
Received path trace: t3-0/3/0:0
  74 33 2d 30 2f 33 2f 30 3a 30 00 00 00 00 0d 0a   t3-0/3/0:0......
Transmitted path trace: t3-0/3/0:0
  74 33 2d 30 2f 33 2f 30 3a 30 00 00 00 00 00 00   t3-0/3/0:0......
Packet Forwarding Engine configuration:
  Destination slot: 0, PLP byte: 1 (0x00)
CoS information:
  CoS transmit queue             Bandwidth           Buffer Priority   Limit
                            %          bps  %          usec
  0 best-effort            95     42499200  95            0       low    none
  3 network-control         5      2236800   5            0       low    none
```

 Copyright © 2018, Juniper Networks, Inc.

```
  Logical interface t3-0/3/0:0.0 (Index 11) (SNMP ifIndex 268) (Generation 499)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 4470, Generation: 578, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 198.51.100.0/24, Local: 198.51.100.1, Broadcast:
198.51.100.255,
 Generation: 98
 DLCI 100
      Flags: Active, Dce-configured
      Total down time: 0 sec, Last down: Never
      Traffic statistics:
       Input  bytes  :                    0
       Output bytes  :                    0
       Input  packets:                    0
       Output packets:                    0
  DLCI statistics:
    Active DLCI  :2 Inactive DLCI  : 0
```

Copyright © 2018, Juniper Networks, Inc.

show interfaces (Channelized OC3 IQ and IQE)

Syntax

show interfaces (*type-fpc/pic/port* <:*channel*><:*channel*><:*channel*>)
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced before Junos OS Release 7.4.

Description

(M Series and T Series routers only) Display status information about the specified channelized OC3 IQ or IQE interface.

Options

type-fpc/pic/port*:*channel*:*channel*:*channel—Interface type with optional corresponding channel levels. The interface type can be one of the following:

- ***type-fpc/pic/port***—For the physical interface, ***type*** is **coc3**. For the clear channel, ***type*** is **so** (for OC3)**.**
- ***type-fpc/pic/port*:*channel***—At the first level of channelization, ***type*** can be **coc1** (channelized OC1), **ct3** (from **coc1)**, or **t3** (from **coc1**).
- ***type-fpc/pic/port*:*channel*:*channel***—At the second level of channelization, ***type*** can be **ct1** (from **coc1** or **ct3)** or **t1** (from **coc1** or **ct3)**.
- ***type-fpc/pic/port*:*channel*:*channel*:*channel***—At the third level of channelization, ***type*** can be **ds** (from **ct1**).

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

List of Sample Output

show interfaces extensive (Channelized OC3 IQ) (Physical) on page 1284
show interfaces extensive (Channelized OC1 on Channelized OC3 IQ) on page 1285
show interfaces extensive (Channelized T1 on Channelized OC3 IQ) on page 1286
show interfaces extensive (DS0 on Channelized OC3 IQ) on page 1287

 Copyright © 2018, Juniper Networks, Inc.

Output Fields Table 62 on page 1271 lists the output fields for the **show interfaces** (all Channelized OC interfaces) command. Output fields are listed in the approximate order in which they appear.

Table 62: Channelized OC show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under “Common Output Fields Description” on page 1110.	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
Description	Interface description.	All levels
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	MTU size on the physical interface.	All levels
Clocking	SONET/SDH reference clock source. It can be **Internal** or **External**. Clocking is configured and displayed only for channel 0.	All levels
***Framing* mode**	Framing mode: **SONET** or **SDH**.	All levels
Speed	Speed at which the interface is running.	All levels
Loopback	Whether loopback is enabled and the type of loopback (**local** or **remote**).	All levels
SONET loopback	Whether loopback is enabled on a SONET/SDH interface, and the type of loopback (**local** or **remote**).	All levels
FCS	Frame check sequence on the interface (either **16** or **32**). The default is **16-bit**.	All levels
Payload scrambler	Whether payload scrambling is enabled.	All levels
Parent	Name and interface index of the interface to which a particular child interface belongs. **None** indicates that this interface is the top level.	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 1110.	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 1110.	All levels

Copyright © 2018, Juniper Networks, Inc.

Table 62: Channelized OC show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Link flags	Information about the link. Possible values are described in the “Link Flags” section under “Common Output Fields Description” on page 1110.	All levels
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
ANSI or ITU LMI settings	(Frame Relay) Settings for Local Management Interface (LMI). The format is **(ANSI or ITU) LMI settings:** ***value***, ***value***... ***nn*** seconds, where ***value*** can be: • **n391dte**—DTE full status polling interval (1–255) • **n392dce**—DCE error threshold (1–10) • **n392dte**—DTE error threshold (1–10) • **n393dce**—DCE monitored event count (1–10) • **n393dte**—DTE monitored event count (1–10) • **t391dte**—DTE polling timer (5–30 seconds) • **t392dce**—DCE polling verification timer (5–30 seconds)	All levels
LMI statistics	(Frame Relay) Statistics about the link management. • **Input**—Number of packets coming in on the interface (*nn*) and how much time has passed since the last packet arrived. The format is **Input:** ***nn*** **(last sent** ***hh:mm:ss*** **ago)**. • **Output**—Number of packets sent out on the interface (*nn*) and how much time has passed since the last packet was sent. The format is **Output:** ***nn*** **(last sent** ***hh:mm:ss*** **ago)**.	**detail extensive**
DTE statistics	(Frame Relay) Statistics about messages transmitted from the data terminal equipment (DTE) to the data circuit-terminating equipment (DCE): • **Enquiries sent**—Number of link status enquiries sent from the DTE to the DCE. • **Full enquiries sent**—Number of full enquiries sent from the DTE to the DCE. • **Enquiry responses received**—Number of enquiry responses received by the DTE from the DCE. • **Full enquiry responses received**—Number of full enquiry responses sent from the DTE to the DCE.	**detail extensive** none
DCE statistics	(Frame Relay) Statistics about messages transmitted from the DCE to the DTE: • **Enquiries received**—Number of enquiries received by the DCE from the DTE. • **Full enquiries received**—Number of full enquiries received by the DCE from the DTE. • **Enquiry responses sent**—Number of enquiry responses sent from the DCE to the DTE. • **Full enquiry responses sent**—Number of full enquiry responses sent from the DCE to the DTE.	**detail extensive** none

Copyright © 2018, Juniper Networks, Inc.

Table 62: Channelized OC show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Common statistics	(Frame Relay) Statistics about messages sent between the DTE and the DCE: • **Unknown messages received**—Number of received packets that do not fall into any category. • **Asynchronous updates received**—Number of link status peer changes received. • **Out-of-sequence packets received**—Number of packets for which the sequence of the packets received is different from the expected sequence. • **Keepalive responses timedout**—Number of keepalive responses that timed out when no LMI packet was reported for **n392dte** or **n393dce** intervals. (See **LMI settings**.)	**detail extensive** none
Nonmatching DCE-end DLCIs	(Frame Relay) Number of DLCIs configured from the DCE, displayed only from the DTE.	**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 hh:mm:ss timezone year-month-day*** **(*hh:mm:ss* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.	**detail extensive** none
CoS Queues	Number of CoS queues configured.	**detail extensive** none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
DS1 alarms **DS1 defects**	E1or T1 media-specific defects that can prevent the interface from passing packets. When a defect persists for a certain period, 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. See the following list for all possible alarms and defects. For complete explanations of most of these alarms and defects, see *Bellcore Telcordia GR-499-CORE*. • **LOS**—Loss of signal. • **LOF**—Loss of frame. • **AIS**—Alarm indication signal. • **YLW**—Yellow alarm. Indicates errors at the remote site receiver.	**detail extensive** none

Table 62: Channelized OC show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
T1 media	Counts of T1 or E1 media-specific errors. • **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. The T1 or E1 media-specific error types are: • **SEF**—Severely errored framing • **BEE**—Bit error • **AIS**—Alarm indication signal • **LOF**—Loss of frame • **LOS**—Loss of signal • **YELLOW**—Errors at the remote site receiver • **BPV**—Bipolar violation • **EXZ**—Excessive zeros • **LCV**—Line code violation • **PCV**—Pulse code violation • **CS**—Carrier state • **FEBE**—Far-end block error (E1 only) • **LES**—Line error seconds • **ES**—Errored seconds • **BES**—Bit error seconds • **SES**—Severely errored seconds • **SEFS**—Severely errored framing seconds • **UAS**—Unavailable seconds	**extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 62: Channelized OC show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Giants**—Number of frames received that are larger than the giant threshold. • **Bucket Drops**—Drops caused by traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off. • **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. • **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. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • **SRAM errors**—Number of hardware errors that occurred in the static RAM (SRAM) on the PIC. If the value of this field increments, the PIC is malfunctioning. • **HS link FIFO overflows**—Number of FIFO overflows on the high-speed links between the ASICs responsible for handling the router interfaces. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 62: Channelized OC show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions** —Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC 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. • **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. • **HS link FIFO underflows**—Number of FIFO underflows on the high-speed links between the ASICs responsible for handling the router interfaces. • **MTU errors**—Number of packets whose size exceeds 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	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**
Active alarms **Active defects**	Defects that can prevent the interface from passing packets: • **None**—There are no active defects or alarms. • **LOF**—Loss of frame.	**detail extensive**
SONET alarms **SONET defects**	Media-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 light the red or yellow alarm LED on the craft interface. See these fields for possible alarms and defects: **SONET PHY**, **SONET section**, **SONET line**, and **SONET path**.	All levels

 Copyright © 2018, Juniper Networks, Inc.

Table 62: Channelized OC show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
SONET vt	SONET virtual-tributary (VT) alarms and defects: • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **BIP-B2**—Bit interleaved parity for SONET line overhead • **REI-V**—Remote error indication (near-end VT) • **LOP-V**—Loss of pointer (near-end VT) • **AIS-V**—Alarm indication signal (near-end VT) • **RDI-V**—Remote defect indication (near-end VT) • **UNEQ-V**—Unequipped (near-end VT) • **PLM-V**—Payload label mismatch (near-end VT) • **ES-V**—Errored seconds (near-end VT) • **SES-V**—Severely errored seconds (near-end VT) • **UAS-V**—Unavailable seconds (near-end VT) • **ES-VFE**—Errored seconds (far-end VT) • **SES-VFE**—Severely errored seconds (far-end VT) • **UAS-VFE**—Unavailable seconds (far-end VT)	**extensive**
SONET PHY	Counts of specific SONET errors with detailed information: • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **PLL Lock**—Phase-locked loop • **PHY Light**—Loss of optical signal	**extensive**
SONET section	Counts of specific SONET errors with detailed information: • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **BIP-B1**—Bit interleaved parity for SONET section overhead • **SEF**—Severely errored framing • **LOS**—Loss of signal • **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**

Table 62: Channelized OC show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
SONET line	Active alarms and defects, plus counts of specific SONET errors with detailed information: • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State** —State of the error. State other than **OK** indicates a problem. Subfields are: • **BIP-B2**—Bit interleaved parity for SONET line overhead • **REI-L**—Remote error indication (near-end line) • **RDI-L**—Remote defect indication (near-end line) • **AIS-L**—Alarm indication signal (near-end line) • **BERR-SF**—Bit error rate fault (signal failure) • **BERR-SD**—Bit error rate defect (signal degradation) • **ES-L**—Errored seconds (near-end line) • **SES-L**—Severely errored seconds (near-end line) • **UAS-L**—Unavailable seconds (near-end line) • **ES-LFE**—Errored seconds (far-end line) • **SES-LFE**—Severely errored seconds (far-end line) • **UAS-LFE**—Unavailable seconds (far-end line)	**extensive**
SONET path	Active alarms and defects, plus counts of specific SONET errors with detailed information: • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **BIP-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) • **ES-PFE**—Errored seconds (far-end STS path) • **SES-PFE**—Severely errored seconds (far-end STS path) • **UAS-PFE**—Unavailable seconds (far-end STS path)	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 62: Channelized OC show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Received SONET overhead **Transmitted SONET overhead**	Values of the received and transmitted SONET/SDH overhead: **F1**—Section user channel byte. This byte is set aside for the purposes of users. **S1**—Synchronization Status (S1). The S1 byte is located in the first STS-1 of an STS-*N*. Bits 5 through 8 convey the synchronization status of the network element. **Z3** and **Z4**—Path overhead. **V5**—Virtual Tributary (VT) path overhead byte.	**extensive**
SDH alarms **SDH defects**	SDH media-specific defects that can prevent the interface from passing packets. When a defect persists for a certain period, 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 light the red or yellow alarm LED on the craft interface. See these fields for possible alarms and defects: SDH PHY, SDH regenerator section, SDH multiplex section, and SDH path. NOTE: For controller based SONET PICs, the SDH alarms and SDH defects output in the **show interface coc3 extensive** command output only shows the section and line level defects. The path level defects can be found under the SONET (so) interface output.	All levels
SDH PHY	Active alarms and defects, plus counts of specific SDH errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **PLL Lock**—Phase-locked loop • **PHY Light**—Loss of optical signal	**extensive**
SDH regenerator section	Active alarms and defects, plus counts of specific SDH errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **RS-BIP8**—24-bit BIP for multiplex section overhead (B2 bytes) • **OOF**—Out of frame • **LOS**—Loss of signal • **LOF**—Loss of frame • **RS-ES**—Errored seconds (near-end regenerator section) • **RS-SES**—Severely errored seconds (near-end regenerator section) • **RS-SEFS**—Severely errored framing seconds (regenerator section)	**extensive**

Table 62: Channelized OC show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
SDH multiplex section	Active alarms and defects, plus counts of specific SDH errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **MS-BIP24**—8-bit BIP for high-order path overhead (B3 byte) • **MS-FEBE**—Far-end block error (multiplex section) • **MS-FERF**—Far-end remote fail (multiplex section) • **MS-AIS**—alarm indication signal (multiplex section) • **BERR-SF**—Bit error rate fault (signal failure) • **BERR-SD**—Bit error rate defect (signal degradation) • **MS-ES**—Errored seconds (near-end multiplex section) • **MS-SES**—Severely errored seconds (near-end multiplex section) • **MS-UAS**—Unavailable seconds (near-end multiplex section) • **MS-ES-FE**—Errored seconds (far-end multiplex section) • **MS-SES-FE**—Severely errored seconds (far-end multiplex section) • **MS-UAS-FE**—Unavailable seconds (far-end multiplex section	**extensive**
SDH path	Active alarms and defects, plus counts of specific SDH errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **HP-BIP8**—8-bit BIP for regenerator section overhead (B1 byte) • **HP-FEBE**—Far-end block error (high-order path) • **HP-LOP**—Loss of pointer (high-order path) • **HP-AIS**—High-order-path alarm indication signal • **HP-FERF**—Far-end remote fail (high-order path) • **HP-UNEQ**—Unequipped (high-order path) • **HP-PLM**—Payload label mismatch (high-order path) • **HP-ES**—Errored seconds (near-end high-order path) • **HP-SES**—Severely errored seconds (near-end high-order path) • **HP-UAS**—Unavailable seconds (near-end high-order path) • **HP-ES-FE**—Errored seconds (far-end high-order path) • **HP-SES-FE**—Severely errored seconds (far-end high-order path) • **HP-UAS-FE**—Unavailable seconds (far-end high-order path)	**extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 62: Channelized OC show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Received SDH overhead **Transmitted SDH overhead**	Values of the received and transmitted SONET overhead: • **C2**—Signal label. Allocated to identify the construction and content of the STS-level SPE and for PDI-P. • **F1**—Section user channel byte. This byte is set aside for the purposes of users. • **K1** and **K2**—These bytes are allocated for APS signaling for the protection of the multiplex section. • **J0**—Section trace. This byte is defined for STS-1 number 1of an STS-*N* signal. Used to transmit a 1-byte fixed-length string or a 16-byte message so that a receiving terminal in a section can verify its continued connection to the intended transmitter. • **S1**—Synchronization status. The S1 byte is located in the first STS-1 of an STS-*N*. • **Z3** and **Z4**—Allocated for future use.	**extensive**
Received path trace **Transmitted path trace**	Channelized OC12 interfaces allow path trace bytes to be sent inband across the SONET/SDH link. 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. This information is specific to each of the 12 channelized OC12 interfaces.	**extensive**
DS3 media	Counts of T3 media-specific errors. For detailed definitions of the T3 (DS-3) error events (BPV, EXZ, LCV, PCV, and CCV) and performance parameters (LES, PES, PSES, CES, CSES, SEFS, and UAS), see RFC 2496. The DS3 or E3 media-specific error types can be: • **PLL Lock**—Phase-locked loop out of lock • **Reframing**—Frame alignment recovery time • **AIS**—Alarm indication signal • **LOF**—Loss of frame • **LOS**—Loss of signal • **IDLE**—Idle code detected • **YELLOW**—Errors at the remote site receiver • **BPV**—Bipolar violation • **EXZ**—Excessive zeros • **LCV**—Line code violation • **PCV**—(DS3 only) Pulse code violation • **CCV**—(DS3 only) C-bit coding violation • **FEBE**—(DS3 only) Far-end block error • **LES**—Line error seconds • **PES**—(DS3 only) P-bit errored seconds • **PSES**—(DS3 only) P-bit errored seconds (section) • **CES**—(DS3 only) C-bit errored seconds • **CSES**—(DS3 only) C-bit severely errored seconds • **SEFS**—Severely errored framing seconds • **UAS**—Unavailable seconds	**extensive**

Table 62: Channelized OC show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
HDLC configuration	Information about the HDLC configuration. • **Policing bucket**—Configured state of the receiving policer. • **Shaping bucket**—Configured state of the transmitting shaper. • **Giant threshold**—Giant threshold programmed into the hardware. • **Runt threshold**—Runt threshold programmed into the hardware. • **Timeslots**—Configured time slots for the interface. • **Line encoding**—Line encoding used. It is always **HDB3**. • **Byte encoding**—(T1 only) Byte encoding used: **Nx64K** or **Nx56K**. • **Line encoding**—Line encoding used. For T1, the value can be **B8ZS** or **AMI**. For E1, the value is **HDB3**. • **Data inversion**—HDLC data inversion setting: **Enabled** or **Disabled**. • **Idle cycle flag**—Idle cycle flags. • **Start end flag**—Start and end flag.	**extensive**
Interface transmit queues	Name of the transmit queues and their associated statistics for each DS3 channel on the Channelized OC12 PIC. • **B/W**—Queue bandwidth as a percentage of the total interface bandwidth. • **WRR**—Weighted round-robin (in percent). • **Packets**—Number of packets transmitted. • **Bytes**—Number of bytes transmitted. • **Drops**—Number of packets dropped. • **Errors**—Number of packet errors.	**extensive**
DSU configuration	Information about the DSU configuration. The last three lines (**Bit count**, **Error bit count**, and **LOS information**) are displayed only if a BERT has ever been run on the interface. • **Compatibility mode**—CSU/DSU compatibility mode: **None**, **Larscom**, **Kentrox**, or **Digital-Link**. • **Scrambling**—Payload scrambling. It can be **Enabled** or **Disabled**. • **Subrate**—Configured subrate setting. Applies only when**Digital-Link** compatibility mode is used. It can be **Disabled** or display units in kbps. • **FEAC loopback**—(T3) Whether a far-end alarm and control (FEAC) loopback is **Active** or **Inactive**. This feature is used to send alarm or status information from the far-end terminal back to the near-end terminal and to initiate T3 loopbacks at the far-end terminal from the near-end terminal. • **Response**—Whether the FEAC signal is **Enabled** or **Disabled**. • **Count**—Number of FEAC loopbacks.	**extensive**
BERT configuration	(DS interfaces) BERT (bit error rate test) checks the quality of the line. This output appears only when a BERT is run on the interface. • **BERT time period**—Configured total time period that the BERT is to run. • **Elapsed**—Actual time elapsed since the start of the BERT (in seconds). • **Induced error rate**—Configured rate at which the bit errors are induced in the BERT pattern. • **Algorithm**—Type of algorithm selected for the BERT.	**detail extensive** none

 Copyright © 2018, Juniper Networks, Inc.

Table 62: Channelized OC show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number. • **PLP byte**—Packet Level Protocol byte.	**extensive**
CoS information	Information about the CoS queue for the physical interface. • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	Logical interface SNMP interface index number.	**detail extensive** none
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under "Common Output Fields Description" on page 1110.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. • **Input rate**—Rate of bits and packets received on the interface. • **Output rate**—Rate of bits and packets transmitted on the interface.	**detail extensive**
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	**detail extensive**

Table 62: Channelized OC show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	**detail extensive**
Protocol	Protocol family configured on the logical interface, such as **iso**, **inet6**, or **mpls**.	**detail extensive** none
Multilink bundle	(If the logical interface is configured as part of a multilink bundle.) Interface name for the multilink bundle.	**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**
Addresses, Flags	Information about the address flags. Possible values are described in the "Addresses Flags" section under "Common Output Fields Description" on page 1110.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address.	**detail extensive** none
DLCI	(Frame Relay) DLCI number of the logical interface. The following DLCI information is displayed: **Flags**, **Total down time**, **Last down**, and **Traffic statistics**. **Flags** is one or more of the following: • **Active**—Set when the link is active and the DTE and DCE are exchanging information. • **Down**—Set when the link is active, but no information is received from the DCE. • **Unconfigured**—Set when the corresponding DLCI in the DCE is not configured. • **Configured**—Set when the corresponding DLCI in the DCE is configured. • **Dce-configured**—Displayed when the command is issued from the DTE.	**detail extensive** none
DLCI statistics	(Frame Relay) Data-link connection identifier (DLCI) statistics. • **Active DLCI**—Number of active DLCIs. • **Inactive DLCI**—Number of inactive DLCIs.	**detail extensive** none

Sample Output

show interfaces extensive (Channelized OC3 IQ) (Physical)

`user@host>` **show interfaces extensive coc3-0/0/0**

Copyright © 2018, Juniper Networks, Inc.

```
Physical interface: coc3-0/0/0, Enabled, Physical link is Down
  Interface index: 128, SNMP ifIndex: 22, Generation: 11
  Description: pink coc3-0/0/0
  Link-level type: Controller, Clocking: Internal, SONET mode, Speed: OC3,
  Loopback: None, Parent: None
  Device flags   : Present Running Down
  Interface flags: Hardware-Down Point-To-Point SNMP-Traps 16384
  Link flags     : None
  Hold-times     : Up 0 ms, Down 0 ms
  CoS queues     : 4 supported
  Last flapped   : 2005-01-27 16:39:21 PST (1w0d 22:09 ago)
  Statistics last cleared: Never
  SONET alarms   : PLL, LOS
  SONET defects  : PLL, LOF, LOS, SEF, AIS-L
  SONET PHY:            Seconds        Count  State
    PLL Lock             681767            1  PLL Lock Error
    PHY Light                 0            0  OK
  SONET section:
    BIP-B1                    0            0
    SEF                  681767            1  Defect Active
    LOS                  681767            1  Defect Active
    LOF                  681767            1  Defect Active
    ES-S                 681767
    SES-S                681767
    SEFS-S               681767
  SONET line:
    BIP-B2                    0            0
    REI-L                     0            0
    RDI-L                     0            0  OK
    AIS-L                681767            1  Defect Active
    BERR-SF                   0            0  OK
    BERR-SD                   0            0  OK
    ES-L                 681767
    SES-L                681767
    UAS-L                681757
    ES-LFE                    0
    SES-LFE                   0
    UAS-LFE                   0
  Received SONET overhead:
    F1      : 0x00, J0      : 0x00, K1      : 0xff, K2      : 0xff
    S1      : 0xff
  Transmitted SONET overhead:
    F1      : 0x00, J0      : 0x01, K1      : 0x00, K2      : 0x00
    S1      : 0x00
```

show interfaces extensive (Channelized OC1 on Channelized OC3 IQ)

```
user@host> show interfaces extensive coc1-0/0/0:1
Physical interface: coc1-0/0/0:1, Enabled, Physical link is Down
  Interface index: 133, SNMP ifIndex: 27, Generation: 16
  Link-level type: Controller, Clocking: Internal, SONET mode, Speed: 51840kbps,

   Loopback: None, Parent: coc3-0/0/0
 Interface index 128
  Device flags   : Present Running Down 16384
  Interface flags: Hardware-Down Point-To-Point SNMP-Traps 16384
  Link flags     : None
  Hold-times     : Up 0 ms, Down 0 ms
  CoS queues     : 4 supported
  Last flapped   : 2005-02-04 14:51:07 PST (00:00:35 ago)
  Statistics last cleared: Never
```

```
  SONET alarms   : None
  SONET defects  : AIS-P
  SONET path:
    BIP-B3                    0           0
    REI-P                     0           0
    LOP-P                     0           0  OK
    AIS-P                    36           1  Defect Active
    RDI-P                     0           0  OK
    UNEQ-P                    0           0  OK
    PLM-P                     0           0  OK
    ES-P                     36
    SES-P                    36
    UAS-P                    26
    ES-PFE                    0
    SES-PFE                   0
    UAS-PFE                   0
  Received SONET overhead:
    C2      : 0xff, C2(cmp) : 0x01, F2      : 0x00, Z3      : 0x00
    Z4      : 0x00, S1(cmp) : 0x00
  Transmitted SONET overhead:
    C2      : 0x01, F2      : 0x00, Z3      : 0x00, Z4      : 0x00
  Received path trace:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
  Transmitted path trace: router-1 coc1-0/0/0:1
    6b 61 76 65 72 69 20 63 6f 63 31 2d 30 2f 30 2f   router-1 coc1-0/0/0:1
    30 3a 31 00 00 00 00 00 00 00 00 00 00 00 00 00   ..............
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
  Packet Forwarding Engine configuration:
    Destination slot: 0 (0x00)
```

show interfaces extensive (Channelized T1 on Channelized OC3 IQ)

```
user@host> show interfaces extensive ct1-0/0/0:1:1
Physical interface: ct1-0/0/0:1:1, Enabled, Physical link is Down
  Interface index: 134, SNMP ifIndex: 62, Generation: 17
  Link-level type: Controller, Clocking: Internal, Speed: T1, Loopback: None,
  Framing: ESF, Parent: coc1-0/0/0:1 Interface index 133
  Device flags   : Present Running Down 16384
  Interface flags: Hardware-Down Point-To-Point SNMP-Traps 16384
  Link flags     : None
  Hold-times     : Up 0 ms, Down 0 ms
  CoS queues     : 4 supported
  Last flapped   : 2005-02-04 14:54:35 PST (00:00:18 ago)
  Statistics last cleared: Never
  DS1   alarms   : None
  DS1   defects  : AIS, LOF
  T1  media:            Seconds        Count  State
    SEF                       1           1  OK
    BEE                       1           1  OK
    AIS                      18           1  Defect Active
    LOF                      18           1  Defect Active
    LOS                       0           0  OK
    YELLOW                    0           0  OK
    BPV                       0           0
    EXZ                       0           0
    LCV                       0           0
    PCV                       0           0
```

 Copyright © 2018, Juniper Networks, Inc.

```
    CS                         0            0
    LES                       18
    ES                        18
    SES                       18
    SEFS                      18
    BES                        0
    UAS                       14
  DS1 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Induced Error rate: 10e-0, Algorithm: 2^15 - 1, 0.151, Pseudorandom (9)
  SONET alarms   : None
  SONET defects  : None
  SONET vt:
    BIP-BIP2                   0            0
    REI-V                      0            0
    LOP-V                      0            0  OK
    AIS-V                     19            1  Defect Active
    RDI-V                     19            1  Defect Active
    UNEQ-V                     0            0  OK
    PLM-V                     19            1  Defect Active
    ES-V                      19
    SES-V                     19
    UAS-V                      9
    ES-VFE                     0
    SES-VFE                    0
    UAS-VFE                    0
  Received SONET overhead:
    V5      : 0x07, V5(cmp) : 0x02
  Transmitted SONET overhead:
    V5      : 0x02
  Packet Forwarding Engine configuration:
    Destination slot: 0 (0x00)
```

show interfaces extensive (DS0 on Channelized OC3 IQ)

```
user@host> show interfaces extensive ds-0/0/0:1:1:1
Physical interface: ds-0/0/0:1:1:1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 63, Generation: 18
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: 320kbps,
  Loopback: None, FCS: 16, Parent: ct1-0/0/0:1:1  Interface index 134
  Device flags   : Present Running
  Interface flags: Hardware-Down Point-To-Point SNMP-Traps 16384
  Link flags     : Keepalives
  Hold-times     : Up 0 ms, Down 0 ms
  CoS queues     : 4 supported
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, L3 incompletes: 0, L2 channel errors: 0,
    L2 mismatch timeouts: 0, HS link CRC errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, Aged packets: 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

HDLC configuration:
  Giant threshold: 1514, Runt threshold: 2
  Timeslots      : 1-5
  Byte encoding: Nx64K, Data inversion: Disabled, Idle cycle flag: flags,
  Start end flag: shared
DS0 BERT configuration:
  BERT time period: 10 seconds, Elapsed: 0 seconds
  Induced Error rate: 10e-0, Algorithm: 2^15 - 1, 0.151, Pseudorandom (9)
Packet Forwarding Engine configuration:
  Destination slot: 0, PLP byte: 4 (0x00)
```

 Copyright © 2018, Juniper Networks, Inc.

show interfaces (Channelized STM1 IQ)

Syntax

show interfaces (*type-fpc/pic /port* <:*channel*><:*channel*>)
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced before Junos OS Release 7.4.

Description

(M Series and T Series routers only) Display status information about the specified channelized STM1 IQ interface.

Options

type-fpc/pic/port*:*channel*:*channel—Interface type with optional corresponding channel levels. The interface type can be one of the following types:

- ***type-fpc/pic/port*:*channel***—For the physical channelized STM1 IQ interface, ***type*** is **cstm1.** For the clear channel, ***type*** is **so**. For channelization, the STM1 IQ interface must be converted to interface type **cau4**.
- ***type-fpc/pic/port*:*channel***—At the first level of channelization, ***type*** can be **ce1** or **e1 (** clear channel or fractional channel from **cau4)**.
- ***type-fpc/pic/port*:*channel*:*channel***—At the second level of channelization, ***type*** is **ds** (from **ce1**).

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

List of Sample Output

show interfaces (Channelized STM1 IQ) (Physical) on page 1290
show interfaces (Channelized AU-4) (Physical) on page 1290
show interfaces (Channelized E1) (Physical) on page 1290
show interfaces (DS) on page 1291

Output Fields

See the output field table for the show interfaces (Channelized STM1) command.

Sample Output

show interfaces (Channelized STM1 IQ) (Physical)

```
user@host> show interfaces cstm1-0/0/0
Physical interface: cstm1-0/0/0, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 35
  Link-level type: Frame-relay, Controller, Clocking: Internal, SDH mode,
  Speed: OC3, Loopback: None, Parent: None  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives DTE
  ANSI LMI settings: n391dte 6, n392dte 3, n393dte 4, t391dte 10 seconds
  LMI: Input: 51700 (00:00:02 ago), Output: 51701 (00:00:02 ago)
  DTE statistics:
    Enquiries sent                      : 43186
    Full enquiries sent                 : 8515
    Enquiry responses received          : 43185
    Full enquiry responses received     : 8515
  DCE statistics:
    Enquiries received                  : 0
    Full enquiries received             : 0
    Enquiry responses sent              : 0
    Full enquiry responses sent         : 0
  Common statistics:
    Unknown messages received           : 0
    Asynchronous updates received       : 0
    Out-of-sequence packets received    : 0
    Keepalive responses timedout        : 0
  Nonmatching DCE-end DLCIs:
          2
  Last flapped   : 2003-02-06 15:01:56 PST (07:15:06 ago)
...
```

show interfaces (Channelized AU-4) (Physical)

```
user@host> show interfaces cau4-0/0/0
Physical interface: cau4-0/0/0, Enabled, Physical link is Up
  Interface index: 147, SNMP ifIndex: 36
  Link-level type: Controller, Clocking: Internal, SDH mode, Speed: OC3,
  Loopback: None, Parent: cstm1-0/0/0 Interface index 146
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : None
  Last flapped   : 2003-02-06 19:36:31 PST (02:40:42 ago)
  SDH   alarms   : None
  SDH   defects  : None
...
```

show interfaces (Channelized E1) (Physical)

```
user@host> show interfaces ce1-0/0/0:11
Physical interface: ce1-0/0/0:11, Enabled, Physical link is Up
  Interface index: 169, SNMP ifIndex: 288
  Link-level type: Frame-relay, Controller, Clocking: Internal, Speed: E1,
  Loopback: None, Framing: G704, Parent: cau4-0/0/0 Interface index 147
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives DTE
  ANSI LMI settings: n391dte 6, n392dte 3, n393dte 4, t391dte 10 seconds
```


Copyright © 2018, Juniper Networks, Inc.

```
  LMI: Input: 51700 (00:00:02 ago), Output: 51701 (00:00:02 ago)
  DTE statistics:
    Enquiries sent                    : 43186
    Full enquiries sent               : 8515
    Enquiry responses received        : 43185
    Full enquiry responses received   : 8515
  DCE statistics:
    Enquiries received                : 0
    Full enquiries received           : 0
    Enquiry responses sent            : 0
    Full enquiry responses sent       : 0
  Common statistics:
    Unknown messages received         : 0
    Asynchronous updates received     : 0
    Out-of-sequence packets received  : 0
    Keepalive responses timedout      : 0
  Nonmatching DCE-end DLCIs:
          2
  Last flapped   : 2003-02-06 22:05:23 PST (00:13:45 ago)
  DS1   alarms   : None
  DS1   defects  : None
  SDH   alarms   : None
  SDH   defects  : None
...
```

show interfaces (DS)

```
user@host> show interfaces ds-0/0/0:11:1
Physical interface: ds-0/0/0:11:1, Enabled, Physical link is Up
  Interface index: 170, SNMP ifIndex: 289
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: 640kbps,
  Loopback: Illegal, FCS: 16, Parent: ce1-0/0/0:11 Interface index 169
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  CoS Queues: 8 maximum usable queues, 4 in use
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 0 (never), Output: 0 (never)
  LCP state: Conf-req-sent
  Egress queues: 8 supported, 4 in use
...
  Logical interface ds-0/0/0:11:1.0 (Index 77) (SNMP ifIndex 290)
    Flags: Hardware-Down Point-To-Point SNMP-Traps Encapsulation: PPP
    Bandwidth: 0
    Protocol inet, MTU: 1500
      Flags: Protocol-Down
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.134.1.0/30, Local: 10.134.1.1
DLCI 100
      Flags: Active, Dce-configured
      Total down time: 0 sec, Last down: Never
      Traffic statistics:
       Input  bytes  :                    0
       Output bytes  :                    0
       Input  packets:                    0
       Output packets:                    0
...
```

Copyright © 2018, Juniper Networks, Inc.

show interfaces (Channelized STM1)

Syntax

show interfaces e1-*fpc/pic/port*:*e1channel*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced before Junos OS Release 7.4.

Description

(M Series and T Series routers only) Display status information about the specified channelized STM1 interface.

Options

e1-*fpc/pic/port*:*e1channel*—Display standard status information about the specified channelized STM1 interface.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

List of Sample Output

show interfaces extensive (Channelized STM1, SDH) on page 1303

Output Fields

Table 63 on page 1292 lists the output fields for the **show interfaces** (all Channelized STM1 interfaces) command. Output fields are listed in the approximate order in which they appear.

Table 63: Channelized STM1 show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under “Common Output Fields Description” on page 1110.	All levels
Interface index	Physical interface's index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none

 Copyright © 2018, Juniper Networks, Inc.

Table 63: Channelized STM1 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**
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**.	All levels
Speed	Speed at which the interface is running.	All levels
Loopback	Whether loopback is enabled and the type of loopback (**local** or **remote**).	All levels
FCS	Frame check sequence on the interface (either **16** or **32**). The default is **16** bits.	All levels
Framing	Physical layer framing format used on the link. It can be **G704**, **G704-NO-CRC4**, or **Unframed**. The default is **G704**.	All levels
Parent	(Channelized STM1 IQ interfaces only) Name and interface index of the interface to which a particular child interface belongs. **None** indicates that this interface is the top level.	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 1110.	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 1110.	All levels
Link flags	Information about the link. Possible values are described in the “Link Flags” section under “Common Output Fields Description” on page 1110.	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 keepalives. • **interval***seconds*—The time in seconds between successive keepalive requests. The range is **10** seconds through **32,767** seconds, with a default of **10** seconds. • **down-count** *number*—The number of keepalive packets a destination must fail to receive before the network takes a link down. The range is**1** through **255**, with a default of **3**. • **up-count** *number*—The number of keepalive packets a destination must receive to change a link’s status from down to up. The range is **1** through **255**, with a default of **1**.	**detail extensive** none

Copyright © 2018, Juniper Networks, Inc.

Table 63: Channelized STM1 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output

Keepalive statistics	(PPP and HDLC) Information about keepalive packets.
	• **Input**—Number of keepalive packets received by PPP.
	  ◦ **(last seen 00:00:00 ago)**—Time since the last keepalive packet was received, in the format*hh:mm:ss*.
	• **Output**—Number of keepalive packets sent by PPP and how long ago the last keepalive packets were sent and received.
	  ◦ **(last seen 00:00:00 ago)**—Time since the last keepalive packet was sent, in the format *hh:mm:ss*.
		detail extensive none

ANSI LMI settingsor **ITU LMI settings**	(Frame Relay) Local Management Interface settings. The format is **(ANSI or ITU) LMI settings:** ***value*****,** ***value*****...** ***xx*** seconds, where ***value*** can be: • **n391dte**—DTE full status polling interval (1- 255) • **n392dce**—DCE error threshold (1-10) • **n392dte**—DTE error threshold (1-10) • **n393dce**—DCE monitored event count (1-10) • **n393dte**—DTE monitored event count (1-10) • **t391dte**—DTE polling timer (5-30 seconds) • **t392dce**—DCE polling verification timer (5-30 seconds)	**detail extensive** none
LMI	(Frame Relay) Statistics about the link management. • **Input**—Number of packets coming in on the interface (*nn*) and how much time has passed since the last packet arrived. The format is **Input:** ***nn*** **(last seen** ***hh*****:*****mm*****:*****ss*** **ago)**. • **Output**—Number of packets sent out on the interface (*nn*) and how much time has passed since the last packet was sent. The format is **Output:** ***nn*** **(last seen** ***hh*****:*****mm*****:*****ss*** **ago)**.	**detail extensive** none
DTE statistics	(Frame Relay) Statistics about messages transmitted from the data terminal equipment (DTE) to the data circuit-terminating equipment (DCE): • **Enquiries sent**—Number of link status enquiries sent from the DTE to the DCE. • **Full enquiries sent**—Number of full enquiries sent from the DTE to the DCE. • **Enquiry responses received**—Number of enquiry responses received by the DTE from the DCE. • **Full enquiry responses received**—Number of full enquiry responses sent from the DTE to the DCE.	**detail extensive** none
DCE statistics	(Frame Relay) Statistics about messages transmitted from the DCE to the DTE: • **Enquiries received**—Number of enquiries received by the DCE from the DTE. • **Full enquiries received**—Number of full enquiries received by the DCE from the DTE. • **Enquiry responses sent**—Number of enquiry responses sent from the DCE to the DTE. • **Full enquiry responses sent**—Number of full enquiry responses sent from the DCE to the DTE.	**detail extensive** none

Copyright © 2018, Juniper Networks, Inc.

Table 63: Channelized STM1 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Common statistics	(Frame Relay) Statistics about messages sent between the DTE and the DCE: • **Unknown messages received**—Number of received packets that do not fall into any category. • **Asynchronous updates received**—Number of link status peer changes received. • **Out-of-sequence packets received**—Number of packets for which the sequence of the packets received is different from the expected sequence. • **Keepalive responses timedout**—Number of keepalive responses that timed out when no LMI packet was reported for **n392dte** or **n393dce** intervals. (See **LMI settings**.)	**detail extensive** none
Nonmatching DCE-end DLCIs	(Frame Relay, displayed only from the DTE) Number of DLCIs configured from the DCE.	**detail extensive** none
LCP state	(PPP) Link Control Protocol state. • **Conf-ack-received**—Acknowledgment was received. • **Conf-ack-sent**—Acknowledgment 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. • **Conf-ack-received**—Acknowledgment was received. • **Conf-ack-sent**—Acknowledgment was sent. • **Conf-req-sent**—Request was sent. • **Down**—NCP negotiation is incomplete (not yet completed or has failed). • **Not-configured**—NCP is not configured on the interface. • **Opened**—NCP negotiation is successful.	**detail extensive** none
CHAP state	(PPP) Displays the state of the Challenge Handshake Authentication Protocol (CHAP) during its transaction. • **Chap-Chal-received**—Challenge was received but response not yet sent. • **Chap-Chal-sent**—Challenge was sent. • **Chap-Resp-received**—Response was received for the challenge sent, but CHAP has not yet moved into the Success state. (Most likely with RADIUS authentication.) • **Chap-Resp-sent**—Response was sent for the challenge received. • **Closed**—CHAP authentication is incomplete. • **Failure**—CHAP authentication failed. • **Not-configured**—CHAP is not configured on the interface. • **Success**—CHAP authentication was successful.	**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

Copyright © 2018, Juniper Networks, Inc.

Table 63: Channelized STM1 show interfaces 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	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **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. • **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. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • **SRAM errors**—Number of hardware errors that occurred in the static RAM (SRAM) on the PIC. If the value of this field increments, the PIC is malfunctioning.	**extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 63: Channelized STM1 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC 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. • **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.	**extensive**
DS1 alarms **DS1 defects**	E1 media-specific defects that can prevent the interface from passing packets. When a defect persists for a certain period, 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. The following lists all possible alarms and defects. For complete explanations of most of these alarms and defects, see *Bellcore Telcordia GR-499-CORE*. • **LOS**—Loss of signal. • **LOF**—Loss of frame. • **AIS**—Alarm indication signal. • **YLW**—Yellow alarm. Indicates errors at the remote site receiver.	**detail extensive** none
SDH alarms **SDH defects**	SDH media-specific defects that can prevent the interface from passing packets. When a defect persists for a certain period, 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 light the red or yellow alarm LED on the craft interface. See these fields for possible alarms and defects: SDH PHY, SDH regenerator section, SDH multiplex section, and SDH path. NOTE: For controller-based SONET PICs, the SDH alarms and SDH defects output in the **show interface cstm1 extensive** command output only shows the section and line level defects. The path level defects can be found under the SONET (so) interface output.	All levels

Copyright © 2018, Juniper Networks, Inc.

Table 63: Channelized STM1 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
E1 media	Active alarms and defects, plus counts of specific E1 errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Error types can be: • **AIS**—Alarm indication signal • **BEE**—Bit error • **BES**—Bit error seconds • **BPV**—Bipolar violation • **CS**—Carrier state • **ES**—Errored seconds • **EXZ**—Excessive zeros • **FEBE**—Far-end block error • **LCV**—Line code violation • **LES**—Line error seconds • **LOF**—Loss of frame • **LOS**—Loss of signal • **PCV**—Pulse code violation • **SEF**—Severely errored framing • **SEFS-S**—Severely errored framing seconds (section) • **SES**—Severely errored seconds • **UAS**—Unavailable seconds • **YELLOW**—Errors at the remote site receiver	**extensive**
Interface transmit queues	Names of the transmit queues and their associated statistics for each E1 channel on the Channelized STM1-to-E1 PIC. • **B/W**—Queue bandwidth as a percentage of the total interface bandwidth. • **WRR**—Weighted round-robin (in percent). • **Packets**—Number of packets transmitted. • **Bytes**—Number of bytes transmitted. • **Drops**—Number of packets dropped. • **Errors**—Number of packet errors.	**extensive**
HDLC configuration	Information about the HDLC configuration. • **Giant threshold**—Giant threshold programmed into the hardware. • **Runt threshold**—Runt threshold programmed into the hardware. • **Timeslots**—Configured time slots for the interface. • **Line encoding**—Line encoding used. It is always **HDB3**.	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 63: Channelized STM1 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
DS1 BERT configuration	BERT (bit error rate test) checks the quality of the line. This output appears only when a BERT is run on the interface. • **BERT time period**—Configured total time period that the BERT is to run. • **Elapsed**—Actual time elapsed since the start of the BERT (in seconds). • **Induced error rate**—Configured rate at which the bit errors are induced in the BERT pattern. • **Algorithm**—Type of algorithm selected for the BERT.	**detail extensive** none
SDH PHY	Active alarms and defects, plus counts of specific SDH errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **PLL Lock**—Phase-locked loop • **PHY Light**—Loss of optical signal	**extensive**
SDH regenerator section	Active alarms and defects, plus counts of specific SDH errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **RS-BIP8**—24-bit BIP for multiplex section overhead (B2 bytes) • **OOF**—Out of frame • **LOS**—Loss of signal • **LOF**—Loss of frame • **RS-ES**—Errored seconds (near-end regenerator section) • **RS-SES**—Severely errored seconds (near-end regenerator section) • **RS-SEFS**—Severely errored framing seconds (regenerator section)	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 63: Channelized STM1 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
SDH multiplex section	Active alarms and defects, plus counts of specific SDH errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **MS-BIP24**—8-bit BIP for high-order path overhead (B3 byte) • **MS-FEBE**—Far-end block error (multiplex section) • **MS-FERF**—Far-end remote fail (multiplex section) • **MS-AIS**—alarm indication signal (multiplex section) • **BERR-SF**—Bit error rate fault (signal failure) • **BERR-SD**—Bit error rate defect (signal degradation) • **MS-ES**—Errored seconds (near-end multiplex section) • **MS-SES**—Severely errored seconds (near-end multiplex section) • **MS-UAS**—Unavailable seconds (near-end multiplex section) • **MS-ES-FE**—Errored seconds (far-end multiplex section) • **MS-SES-FE**—Severely errored seconds (far-end multiplex section) • **MS-UAS-FE**—Unavailable seconds (far-end multiplex section)	**extensive**
SDH path	Active alarms and defects, plus counts of specific SDH errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **HP-BIP8**—8-bit BIP for regenerator section overhead (B1 byte) • **HP-FEBE**—Far-end block error (high-order path) • **HP-LOP**—Loss of pointer (high-order path) • **HP-AIS**—High-order-path alarm indication signal • **HP-FERF**—Far-end remote fail (high-order path) • **HP-UNEQ**—Unequipped (high-order path) • **HP-PLM**—Payload label mismatch (high-order path) • **HP-ES**—Errored seconds (near-end high-order path) • **HP-SES**—Severely errored seconds (near-end high-order path) • **HP-UAS**—Unavailable seconds (near-end high-order path) • **HP-ES-FE**—Errored seconds (far-end high-order path) • **HP-SES-FE**—Severely errored seconds (far-end high-order path) • **HP-UAS-FE**—Unavailable seconds (far-end high-order path)	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 63: Channelized STM1 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
SDH tu	Active alarms and defects, plus counts of specific SDH tributary unit (TU) errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **TU-BIP-2**—Bit interleaved parity for SONET line overhead • **TU-FEBE**—(near-end TU) • **TU-LOP**—Loss of pointer (near-end TU) • **TU-AIS**—Alarm indication signal (near-end TU) • **TU-FERF**—(near-end TU) • **TU-UNEQ**—Unequipped (near-end TU) • **TU-PLM**—Payload label mismatch (near-end TU) • **TU-ES**—Errored seconds (near-end TU) • **TU-SES**—Severely errored seconds (near-end TU) • **TU-UAS**—Unavailable seconds (near-end TU) • **TU-ES-FE**—Errored seconds (far-end TU) • **TU-SES-FE**—Severely errored seconds (far-end TU) • **TU-UAS-FE**—Unavailable seconds (far-end TU)	**extensive**
Received SDH overhead **Transmitted SDH overhead**	Values of the received and transmitted SONET overhead: • **C2**—Signal label. Allocated to identify the construction and content of the STS-level SPE and for PDI-P. • **F1**—Section user channel byte. This byte is set aside for the purposes of users. • **K1** and **K2**—These bytes are allocated for APS signaling for the protection of the multiplex section. • **J0**—Section trace. This byte is defined for STS-1 number 1 of an STS-*N* signal. Used to transmit a 1-byte fixed-length string or a 16-byte message so that a receiving terminal in a section can verify its continued connection to the intended transmitter. • **S1**—Synchronization status. The S1 byte is located in the first STS-1 of an STS-*N* signal. • **Z3** and **Z4**—Allocated for future use.	**extensive**
Received path trace **Transmitted path trace**	Channelized OC12 interfaces allow path trace bytes to be sent inband across the SONET/SDH link. 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. This information is specific to each of the 12 channelized OC12 interfaces.	**extensive**
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number. • **PLP byte**—Packet Level Protocol byte.	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 63: Channelized STM1 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface. • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under “Common Output Fields Description” on page 1110.	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
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under “Common Output Fields Description” on page 1110.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family configured on the logical interface, such as **iso**, **inet6**, or **mpls**.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route table	Routing table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive**
Flags	Information about the protocol family flags. Possible values are described in the “Family Flags” section under “Common Output Fields Description” on page 1110.	**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 1110.	**detail extensive** none

Copyright © 2018, Juniper Networks, Inc.

Table 63: Channelized STM1 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address.	**detail extensive** none
DLCI	(Frame Relay) DLCI number of the logical interface. The following DLCI information is displayed: **Flags**, **Total down time**, **Last down**, and **Traffic statistics**. **Flags** is one or more of the following: • **Active**—Set when the link is active and the DTE and DCE are exchanging information. • **Down**—Set when the link is active, but no information is received from the DCE. • **Unconfigured**—Set when the corresponding DLCI in the DCE is not configured. • **Configured**—Set when the corresponding DLCI in the DCE is configured. • **Dce-configured**—Displayed when the command is issued from the DTE.	**detail extensive** none
DLCI statistics	(Frame Relay) Data-link connection identifier (DLCI) statistics. • **Active DLCI**—Number of active DLCIs. • **Inactive DLCI**—Number of inactive DLCIs.	**detail extensive** none

Sample Output

show interfaces extensive (Channelized STM1, SDH)

```
user@host> show interfaces e1-1/0/0:1 extensive
Physical interface: e1-1/0/0:1, Enabled, Physical link is Up
  Interface index: 148, SNMP ifIndex: 285, Generation: 2915
  Link-level type: Frame-relay, MTU: 1504, SDH mode, Speed: E1, Loopback: None,
  FCS: 16, Framing: G704
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives DTE
  ANSI LMI settings: n391dte 6, n392dte 3, n393dte 4, t391dte 10 seconds
  LMI: Input: 51700 (00:00:02 ago), Output: 51701 (00:00:02 ago)
  DTE statistics:
    Enquiries sent                     : 43186
    Full enquiries sent                : 8515
    Enquiry responses received         : 43185
    Full enquiry responses received    : 8515
  DCE statistics:
    Enquiries received                 : 0
    Full enquiries received            : 0
    Enquiry responses sent             : 0
    Full enquiry responses sent        : 0
  Common statistics:
    Unknown messages received          : 0
    Asynchronous updates received      : 0
    Out-of-sequence packets received   : 0
    Keepalive responses timedout       : 0
```

```
Nonmatching DCE-end DLCIs:
        2
Hold-times     : Up 0 ms, Down 0 ms
Last flapped   : 2002-05-23 17:02:59 PDT (17:23:45 ago)
Statistics last cleared: Never
Traffic statistics:
 Input  bytes  :                 592                  48 bps
 Output bytes  :                 644                  48 bps
 Input  packets:                  46                   0 pps
 Output packets:                  46                   0 pps
Input errors:
  Errors: 0, Drops: 9, Framing errors: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 11, L2 mismatch timeouts: 0,
  HS link CRC errors: 0, SRAM errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, Aged packets: 0
DS1   alarms   : None
DS1   defects  : None
SDH   alarms   : None
SDH   defects  : None
E1  media:            Seconds        Count  State
  SEF                       0            0  OK
  BEE                       0            0  OK
  AIS                     124            1  OK
  LOF                     124            1  OK
  LOS                       0            0  OK
  YELLOW                    0            0  OK
  BPV                       0            0
  EXZ                       0            0
  LCV                       0            0
  PCV                       0            0
  CS                        0            0
  FEBE                      0            0
  LES                     124
  ES                      125
  SES                     124
  SEFS                    124
  BES                       0
  UAS                      37
Interface transmit queues:
            B/W  WRR       Packets        Bytes        Drops       Errors
  Queue0     95   95             0            0            0            0
  Queue1      5    5           529         6348            0            0
HDLC configuration:
  Giant threshold: 0, Runt threshold: 0
  Timeslots      : All active
  Line encoding: HDB3
DS1 BERT configuration:
  BERT time period: 10 seconds, Elapsed: 0 seconds
  Induced Error rate: 10e-0, Algorithm: 2^15 - 1, 0.151, Pseudorandom (9)
SDH PHY:              Seconds        Count  State
  PLL Lock                  0            0  OK
  PHY Light                 0            0  OK
SDH regenerator section:
  RS-BIP8                   0            0
  OOF                     125            1  OK
  LOS                     125            1  OK
  LOF                     125            1  OK
  RS-ES                   125
  RS-SES                  125
  RS-SEFS                 125
```

 Copyright © 2018, Juniper Networks, Inc.

```
SDH multiplex section:
  MS-BIP24                 0            0
  MS-FEBE                  0            0
  MS-FERF                  0            0  OK
  MS-AIS                 125            1  OK
  BERR-SF                  0            0  OK
  BERR-SD                  0            0  OK
  MS-ES                  125
  MS-SES                 125
  MS-UAS                 115
  MS-ES-FE                 0
  MS-SES-FE                0
  MS-UAS-FE                0
SDH path:
  HP-BIP8                  0            0
  HP-FEBE                  0            0
  HP-LOP                   0            0  OK
  HP-AIS                 125            1  OK
  HP-FERF                  0            0  OK
  HP-UNEQ                  0            0  OK
  HP-PLM                 125            1  OK
  HP-ES                  125
  HP-SES                 125
  HP-UAS                 115
  HP-ES-FE                 0
  HP-SES-FE                0
  HP-UAS-FE                0
SDH tu:
  TU-BIP2                  0            0
  TU-FEBE                124            1
  TU-LOP                   0            0  OK
  TU-AIS                 124            1  OK
  TU-FERF                124            1  OK
  TU-UNEQ                  0            0  OK
  TU-PLM                 124            1  OK
  TU-ES                  125
  TU-SES                 125
  TU-UAS                 115
  TU-ES-FE                 0
  TU-SES-FE                0
  TU-UAS-FE                0
Received SDH overhead:
  F1      : 0x00, J0      : 0x00, K1      : 0x00, K2      : 0x00
  S1      : 0x00, C2      : 0x02, C2(cmp) : 0x02, F2      : 0x00
  Z3      : 0x00, Z4      : 0x00, S1(cmp) : 0x00, V5      : 0x02
  V5(cmp) : 0x02
Transmitted SDH overhead:
  F1      : 0x00, J0      : 0x00, K1      : 0x00, K2      : 0x00
  S1      : 0x00, C2      : 0x02, F2      : 0x00, Z3      : 0x00
  Z4      : 0x00, V5      : 0x02
Received path trace:
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
Transmitted path trace:
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
Packet Forwarding Engine configuration:
  Destination slot: 1, PLP byte: 2 (0x07)
CoS information:
  CoS transmit queue            Bandwidth           Buffer Priority   Limit
                           %          bps   %          usec
  0 best-effort           95      1945600  95             0      low    none
  3 network-control        5       102400   5             0      low    none
```

```
  Logical interface e1-1/0/0:1.0 (Index 10) (SNMP ifIndex 369) (Generation 496)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 1500, Generation: 575, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 198.51.100.0/24, Local: 198.51.100.1, Broadcast:
198.51.100.255,
        Generation: 975
    DLCI 100
      Flags: Active, Dce-configured
      Total down time: 0 sec, Last down: Never
      Traffic statistics:
       Input  bytes  :                    0
       Output bytes  :                    0
       Input  packets:                    0
       Output packets:                    0
  DLCI statistics:
    Active DLCI  :2 Inactive DLCI  : 0
```

 Copyright © 2018, Juniper Networks, Inc.

show interfaces (Channelized T1 IQ)

Syntax

show interfaces (ct1-*fpc/pic/port* | *type-fpc/pic/port*<:*channel*><:*channel*>)
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced in Junos OS Release 7.4.

Description

(M Series and T Series routers only) Display status information about the specified channelized T1 IQ interface.

Options

type-fpc/pic/port*:*channel—Interface type. With optional corresponding channel levels, the interface type can be one of the following:

- ***type-fpc/pic/port***—For the physical channelized T1 IQ interface, ***type*** is **ct1.**
- ***type-fpc/pic/port*:*channel***—For the clear channel, ***type*** is **t1**. At the first level of channelization, ***type*** can be **ct1** or **t1.**
- ***type-fpc/pic/port*:*channel*:*channel***—At the second level of channelization, ***type*** can be **ds**.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

List of Sample Output

show interfaces extensive (CT1) on page 1315
show interfaces extensive (T1) on page 1316
show interfaces extensive (DS0) on page 1317

Output Fields

Table 64 on page 1308 lists the output fields for the **show interfaces** (Channelized T1 IQ and T3 IQ interfaces) command. Output fields are listed in the approximate order in which they appear.

Copyright © 2018, Juniper Networks, Inc.

Table 64: Channelized T1 IQ and T3 IQ show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under “Common Output Fields Description” on page 1110.	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**
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**.	All levels
Speed	Speed at which the interface is running.	All levels
Loopback	Whether loopback is enabled and the type of loopback (**local** or **remote**).	All levels
FCS	Frame check sequence on the interface (either **16** or **32**). The default is **16** bits.	All levels
Framing	Physical layer framing format used on the link. It can be **ESF** or **SF**. The default is **ESF**.	All levels
Parent	Name and interface index of the interface to which a particular child interface belongs. **None** indicates that this interface is the top level.	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 1110.	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 1110.	All levels
Link flags	Information about the link. Possible values are described in the “Link Flags” section under “Common Output Fields Description” on page 1110.	All levels
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 64: Channelized T1 IQ and T3 IQ show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Keepalive settings	Configured settings for keepalives. • **interval** ***seconds***—The time in seconds between successive keepalive requests. The range is **10** seconds through **32,767** seconds, with a default of **10** seconds. • **down-count** ***number***—The number of keepalive packets a destination must fail to receive before the network takes a link down. The range is **1** through **255**, with a default of **3**. • **up-count** ***number***—The number of keepalive packets a destination must receive to change a link's status from down to up. The range is **1** through **255**, with a default of **1**.	**detail extensive** none

Keepalive statistics	Information about keepalive packets.
	• **Input**—Number of keepalive packets received by PPP.
	  ◦ **(last seen 00:00:00 ago)**—Time since the last keepalive packet was received, in the format ***hh:mm:ss***.
	• **Output**—Number of keepalive packets sent by PPP and how long ago the last keepalive packets were sent and received.
	  ◦ **(last seen 00:00:00 ago)**—Time since the last keepalive packet was sent, in the format ***hh:mm:ss***.
		detail extensive none

LMI settings	(Frame Relay) Settings for Local Management Interface (LMI) can be either **ANSI LMI settings** or **ITU LMI settings**. ANSI LMI settings is the default. The format is **(ANSI or ITU) LMI settings:** ***value*****,** ***value...*** ***xx seconds***, where ***value*** can be: • **n391dte**—DTE full status polling interval (1–255) • **n392dce**—DCE error threshold (1–10) • **n392dte**—DTE error threshold (1–10) • **n393dce**—DCE monitored event count (1–10) • **n393dte**—DTE monitored event count (1–10) • **t391dte**—DTE polling timer (5–30 seconds) • **t392dce**—DCE polling verification timer (5–30 seconds)	**detail extensive** none
LMI	(Frame Relay) LMI packet statistics: • **Input**—Number of packets coming in on the interface (***nn***) and how much time has passed since the last packet arrived. The format is **Input:** ***nn*** **(last seen** ***hh:mm:ss*** **ago)**. • **Output**—Number of packets sent out on the interface (***nn***) and how much time has passed since the last packet was sent. The format is **Output:** ***nn*** **(last sent** ***hh:mm:ss*** **ago)**.	**detail extensive** none
DTE statistics	(Frame Relay) Statistics about messages transmitted from the data terminal equipment (DTE) to the data communication equipment (DCE): • **Enquiries sent**—Number of link status enquiries sent from the DTE to the DCE. • **Full enquiries sent**—Number of full enquiries sent from the DTE to the DCE. • **Enquiry responses received**—Number of enquiry responses received by the DTE from the DCE. • **Full enquiry responses received**—Number of full enquiry responses sent from the DTE to the DCE.	**detail extensive** none

Copyright © 2018, Juniper Networks, Inc.

Table 64: Channelized T1 IQ and T3 IQ show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
DCE statistics	(Frame Relay) Statistics about messages transmitted from the DCE to the DTE: • **Enquiries received**—Number of enquiries received by the DCE from the DTE. • **Full enquiries received**—Number of full enquiries received by the DCE from the DTE. • **Enquiry responses sent**—Number of enquiry responses sent from the DCE to the DTE. • **Full enquiry responses sent**—Number of full enquiry responses sent from the DCE to the DTE.	**detail extensive** none
Common statistics	(Frame Relay) Statistics about messages sent between the DTE and the DCE: • **Unknown messages received**—Number of received packets that do not fall into any category. • **Asynchronous updates received**—Number of link status peer changes received. • **Out-of-sequence packets received**—Number of packets for which the sequence of the packets received is different from the expected sequence. • **Keepalive responses timedout**—Number of keepalive responses that timed out when no LMI packet was reported for **n392dte** or **n393dce** intervals. (See **LMI settings**.)	**detail extensive** none
Nonmatching DCE-end DLCIs	(Frame Relay) Number of DLCIs configured from the DCE, displayed only from the DTE.	**detail extensive** none
LCP state	(PPP) Link Control Protocol state. • **Conf-ack-received**—Acknowledgement was received. • **Conf-ack-sent**—Acknowledgement was sent. • **Conf-req-sent**—Request was sent. • **Down**—LCP negotiation is incomplete (not yet completed or has failed). • **Not-configured**—LCP is not configured on the interface. • **Opened**—LCP negotiation is successful.	**detail extensive** none
NCP state	(PPP) Network Control Protocol state. • **Conf-ack-received**—Acknowledgement was received. • **Conf-ack-sent**—Acknowledgement was sent. • **Conf-req-sent**—Request was sent. • **Down**—NCP negotiation is incomplete (not yet completed or has failed). • **Not-configured**—NCP is not configured on the interface. • **Opened**—NCP negotiation is successful.	**detail extensive** none

 Copyright © 2018, Juniper Networks, Inc.

Table 64: Channelized T1 IQ and T3 IQ show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
CHAP state	(PPP) Displays the state of the Challenge Handshake Authentication Protocol (CHAP) during its transaction. • **Chap-Chal-received**—Challenge was received but response not yet sent. • **Chap-Chal-sent**—Challenge was sent. • **Chap-Resp-received**—Response was received for the challenge sent, but CHAP has not yet moved into the**Success** state. (Most likely with RADIUS authentication.) • **Chap-Resp-sent**—Response was sent for the challenge received. • **Closed**—CHAP authentication is incomplete. • **Failure**—CHAP authentication failed. • **Not-configured**—CHAP is not configured on the interface. • **Success**—CHAP authentication was successful.	**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
CoS queues	Number of CoS queues configured.	**detail extensive** none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes, 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**

Table 64: Channelized T1 IQ and T3 IQ show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Giants**—Number of frames received that are larger than the giant threshold. • **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. • **L2 channel errors**—Counter increments when the software could not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Count of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **HS link CRC errors**—Count of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • **SRAM errors**—Number of hardware errors that occurred in the static RAM (SRAM) on the PIC. If the value in this field increments, the PIC is malfunctioning. • **Resource errors**—Sum of transmit drops.	**extensive**
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly, (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **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. • **MTU errors**—Number of packets whose size exceeds the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 64: Channelized T1 IQ and T3 IQ show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Queue counters	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**
DS1 alarms DS1 defects	Media-specific defects that can render the interface unable to pass 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. • **LOS**—Loss of signal. • **LOF**—Loss of frame. • **AIS**—Alarm indication signal. • **YLW**—Yellow alarm. Indicates errors at the remote site receiver.	**detail extensive** none
T1 media	Counts of T1 media-specific errors. • **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. The T1 media-specific error types can be: • **AIS**—Alarm indication signal • **BEE**—Bit error event • **BES**—Bit error seconds • **BPV**—Bipolar violation • **CS**—Carrier state • **ES**—Errored seconds • **EXZ**—Excessive zeros • **FEBE**—Far-end block error • **LCV**—Line code violation • **LES**—Line error seconds • **LOF**—Loss of frame • **LOS**—Loss of signal • **PCV**—Pulse code violation • **SEF**—Severely errored framing • **SEFS**—Severely errored framing seconds (section) • **SES**—Severely errored seconds • **UAS**—Unavailable seconds • **YELLOW**—Errors at the remote site receiver	**extensive**
Line encoding	Line encoding used: **B8ZS** or **AMI**.	All levels
Buildout	Buildout setting.	All levels

Copyright © 2018, Juniper Networks, Inc.

Table 64: Channelized T1 IQ and T3 IQ show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
HDLC configuration	Information about the HDLC configuration. • **Policing bucket**—Configured state of the receiving policer. • **Shaping bucket**—Configured state of the transmitting shaper. • **Giant threshold**—Giant threshold programmed into the hardware. • **Runt threshold**—Runt threshold programmed into the hardware. • **Timeslots**—Configured time slots for the interface. • **Line encoding**—Line encoding used: **B8ZS** or **AMI**. • **Byte encoding**—Byte encoding used: **Nx64K** or **Nx56K**. • **Data inversion**—HDLC data inversion setting: **Enabled** or **Disabled**. • **Idle cycle Flag**—Idle cycle flags. • **Start end Flag**—Start and end flag.	**extensive**
DS0 or DS1 BERT configuration	BERT (bit error rate test) checks the quality of the line. This output appears only when a BERT is run on the interface. • **BERT time period**—Configured total time period that the BERT is to run. • **Elapsed**—Actual time elapsed since the start of the BERT (in seconds). • **Induced error rate**—Configured rate at which the bit errors are induced in the BERT pattern. • **Algorithm**—Type of algorithm selected for the BERT.	**detail extensive** none
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number. • **PLP byte**—Packet Level Protocol byte.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	Logical interface SNMP interface index number.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface; values are described in the “Logical Interface Flags” section under “Common Output Fields Description” on page 1110.	**All levels**
Encapsulation	Encapsulation on the logical interface.	**All levels**
Protocol	Protocol family configured on the logical interface, such as **iso**, **inet6**, or **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**

 Copyright © 2018, Juniper Networks, Inc.

Table 64: Channelized T1 IQ and T3 IQ show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Route table	Routing table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive**
Flags	Information about the protocol family flags. Possible values are described in the “Family Flags” section under “Common Output Fields Description” on page 1110.	**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 1110.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
DLCI	(Frame Relay) DLCI number of the logical interface. The following DLCI information is displayed: **Flags**, **Total down time**, **Last down**, and **Traffic statistics**. **Flags** is one or more of the following: • **Active**—Set when the link is active and the DTE and DCE are exchanging information. • **Down**—Set when the link is active, but no information is received from the DCE. • **Unconfigured**—Set when the corresponding DLCI in the DCE is not configured. • **Configured**—Set when the corresponding DLCI in the DCE is configured. • **Dce-configured**—Displayed when the command is issued from the DTE.	**detail extensive** none
DLCI statistics	(Frame Relay) Data-link connection identifier (DLCI) statistics. • **Active DLCI**—Number of active DLCIs. • **Inactive DLCI**—Number of inactive DLCIs.	**detail extensive** none

Sample Output

show interfaces extensive (CT1)

```
user@host> show interfaces extensive ct1-0/1/1
Physical interface: ct1-0/1/1, Enabled, Physical link is Up
   Interface index: 145, SNMP ifIndex: 32, Generation: 28
   Link-level type: Controller, Clocking: Internal, Speed: T1,
   Loopback: None, Framing: ESF, Parent: None
 Device flags   : Present Running
   Interface flags: Point-To-Point SNMP-Traps 16384
   Link flags     : None
   Hold-times     : Up 0 ms, Down 0 ms
   CoS queues     : 4 supported
   Last flapped   : 2005-08-17 11:47:09 PDT (1d 03:38 ago)
```

Copyright © 2018, Juniper Networks, Inc.

```
    Statistics last cleared: 2005-08-18 15:25:37 PDT (00:00:27 ago)
    DS1   alarms   : None
    DS1   defects  : None
    T1  media:            Seconds        Count  State
      SEF                       0            0  OK
      BEE                       0            0  OK
      AIS                       0            0  OK
      LOF                       0            0  OK
      LOS                       0            0  OK
      YELLOW                    0            0  OK
      BPV                       0            0
      EXZ                       0            0
      LCV                       0            0
      PCV                       0            0
      CS                        0            0
      LES                       0
      ES                        0
      SES                       0
      SEFS                      0
      BES                       0
      UAS                       0
      Line encoding: B8ZS
      Buildout       : 0 to 132 feet
    DS1 BERT configuration:
      BERT time period: 10 seconds, Elapsed: 0 seconds
      Induced Error rate: 10e-0, Algorithm: 2^15 - 1, 0.151, Pseudorandom (9)
    Packet Forwarding Engine configuration:
      Destination slot: 0 (0x00)
```

show interfaces extensive (T1)

```
user@host> show interfaces extensive t1-0/2/0
Physical interface: t1-0/2/0, Enabled, Physical link is Up
   Interface index: 161, SNMP ifIndex: 33, Generation: 61
   Link-level type: PPP, MTU: 1504, Speed: T1, Loopback: None, FCS: 16,
   Parent: ct1-0/2/0 Interface index 148
   Device flags   : Present Running
   Interface flags: Point-To-Point SNMP-Traps 16384
   Link flags     : Keepalives DTE
  ANSI LMI settings: n391dte 6, n392dte 3, n393dte 4, t391dte 10 seconds
  LMI: Input: 51700 (00:00:02 ago), Output: 51701 (00:00:02 ago)
  DTE statistics:
    Enquiries sent                     : 43186
    Full enquiries sent                : 8515
    Enquiry responses received         : 43185
    Full enquiry responses received    : 8515
  DCE statistics:
    Enquiries received                 : 0
    Full enquiries received            : 0
    Enquiry responses sent             : 0
    Full enquiry responses sent        : 0
  Common statistics:
    Unknown messages received          : 0
    Asynchronous updates received      : 0
    Out-of-sequence packets received   : 0
    Keepalive responses timedout       : 0
  Nonmatching DCE-end DLCIs:
          2
   Hold-times     : Up 0 ms, Down 0 ms
   CoS queues     : 4 supported
   Last flapped   : 2005-09-07 15:43:47 PDT (00:00:06 ago)
```

 Copyright © 2018, Juniper Networks, Inc.

```
   Statistics last cleared: Never
   Traffic statistics:
    Input  bytes  :                    0                    0 bps
    Output bytes  :                   14                    0 bps
    Input  packets:                    0                    0 pps
    Output packets:                    1                    0 pps
   Input errors:
     Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
     Policed discards: 0, L3 incompletes: 0,
     L2 channel errors: 0, L2 mismatch timeouts: 0,
     HS link CRC errors: 0, SRAM errors: 0, Resource errors: 0
   Output errors:
     Carrier transitions: 1, Errors: 0, Drops: 0, Aged packets: 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                 1                    1               0
   DS1   alarms   : None
   DS1   defects  : None
   HDLC configuration:
     Policing bucket: Disabled
     Shaping bucket : Disabled
     Giant threshold: 1514, Runt threshold: 2
     Timeslots      : All active
     Byte encoding: Nx64K, Data inversion: Disabled, Idle cycle flag:
     flags, Start end flag: shared
   Packet Forwarding Engine configuration:
     Destination slot: 0, PLP byte: 4 (0x00)
```

show interfaces extensive (DS0)

```
user@host> show interfaces extensive ds-0/1/0:0
Physical interface: ds-0/1/0:1, Enabled, Physical link is Up
   Interface index: 157, SNMP ifIndex: 52, Generation: 46
   Link-level type: Frame-Relay, PPP, MTU: 1504, Clocking: Internal,
   Speed: 640kbps, Loopback: None, FCS:16,
   Parent: ct1-0/1/0 Interface index 143
   Device flags   : Present Running
   Interface flags: Point-To-Point SNMP-Traps 16384
   Link flags     : Keepalives DTE
  ANSI LMI settings: n391dte 6, n392dte 3, n393dte 4, t391dte 10 seconds
  LMI: Input: 51700 (00:00:02 ago), Output: 51701 (00:00:02 ago)
  DTE statistics:
    Enquiries sent                    : 43186
    Full enquiries sent               : 8515
    Enquiry responses received        : 43185
    Full enquiry responses received   : 8515
  DCE statistics:
    Enquiries received                : 0
    Full enquiries received           : 0
    Enquiry responses sent            : 0
    Full enquiry responses sent       : 0
  Common statistics:
    Unknown messages received         : 0
    Asynchronous updates received     : 0
    Out-of-sequence packets received  : 0
    Keepalive responses timedout      : 0
  Nonmatching DCE-end DLCIs:
          2
```

```
   Hold-times     : Up 0 ms, Down 0 ms
   Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
   Keepalive statistics:
     Input : 8 (last seen 00:00:12 ago)
     Output: 8 (last sent 00:00:07 ago)
   LCP state: Opened
   NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured,
   mpls: Not-configured
   CHAP state: Not-configured
   CoS queues     : 4 supported
   Last flapped   : 2005-08-18 15:23:46 PDT (00:03:17 ago)
   Statistics last cleared: 2005-08-18 15:25:37 PDT (00:01:26 ago)
   Traffic statistics:
    Input  bytes  :                 840                   0 bps
    Output bytes  :                 912                   0 bps
    Input  packets:                  25                   0 pps
    Output packets:                  26                   0 pps
   Input errors:
     Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
     Policed discards: 0, L3 incompletes: 0,
     L2 channel errors: 0, L2 mismatch timeouts: 0, HS link CRC errors: 0,
     Resource errors: 0
   Output errors:
     Carrier transitions: 0, Errors: 0, Drops: 0, Aged packets: 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              26                   26               0
   HDLC configuration:
     Giant threshold: 1514, Runt threshold: 2
     Timeslots      : 1-10
     Byte encoding: Nx64K, Data inversion: Disabled, Idle cycle flag: flags,
     Start end flag: shared
   DS0 BERT configuration:
     BERT time period: 10 seconds, Elapsed: 0 seconds
     Induced Error rate: 10e-0, Algorithm: 2^15 - 1, 0.151, Pseudorandom (9)
   Packet Forwarding Engine configuration:
     Destination slot: 0, PLP byte: 4 (0x00)
Logical interface ds-0/1/0:1.0 (Index 67) (SNMP ifIndex 53) (Generation 11)
     Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
     Protocol inet, MTU: 1500, Generation: 26, Route table: 0
       Flags: None
       Addresses, Flags: Is-Preferred Is-Primary
         Destination: 11.11.11.0/30, Local: 11.11.11.2, Broadcast: 11.11.11.3,
         Generation: 39
    DLCI 100
      Flags: Active, Dce-configured
      Total down time: 0 sec, Last down: Never
      Traffic statistics:
       Input  bytes  :                   0
       Output bytes  :                   0
       Input  packets:                   0
       Output packets:                   0
  DLCI statistics:
    Active DLCI  :2 Inactive DLCI  : 0
...
```

 Copyright © 2018, Juniper Networks, Inc.

show interfaces (Channelized T3 IQ)

Syntax

show interfaces (ct3-*fpc/pic/port* | *type-fpc/pic/port*<:*channel*><:*channel*>)
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced before Junos OS Release 7.4.

Description

(M Series and T Series routers only) Display status information about the specified channelized T3 IQ interface.

Options

type-fpc/pic/port*:*channel—Interface type. With optional corresponding channel levels, the interface type can be one of the following:

- ***type-fpc/pic/port***—For the physical channelized T3 IQ interface, ***type*** is **ct3**.
- ***type-fpc/pic/port*:*channel***—For the clear channel, ***type*** is **t3**. At the first level of channelization, ***type*** can be **ct1** or **t1**.
- ***type-fpc/pic/port*:*channel*:*channel***—At the second level of channelization, ***type*** is **ds**.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

List of Sample Output

show interfaces extensive (Channelized T3 IQ) (Physical) on page 1320
show interfaces extensive (Channelized T1 on Channelized T3 IQ) on page 1320
show interfaces extensive (DS0 on Channelized T3 IQ) on page 1320

Output Fields

See the output field table for the show interfaces (Channelized T1 IQ) command.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show interfaces extensive (Channelized T3 IQ) (Physical)

```
user@host> show interfaces extensive ct3-0/0/1
Physical interface: ct3-0/0/1, Enabled, Physical link is Up
  Interface index: 30, SNMP ifIndex: 317, Generation: 29
  Link-level type: Controller, MTU: 4474, Clocking: Internal, Speed: T3,
  Loopback: None, FCS: 16, Mode: C/Bit parity, Parent: None
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : None
...
```

show interfaces extensive (Channelized T1 on Channelized T3 IQ)

```
user@host> show interfaces extensive ct1-0/0/1:2
Physical interface: ct1-0/0/1:2, Enabled, Physical link is Up
  Interface index: 175, SNMP ifIndex: 1505, Generation: 174
  Link-level type: Controller, MTU: 1504, Clocking: Internal, Speed: T1,
  Loopback: None, FCS: 16, Framing: ESF, Parent: ct3-0/0/1 (Index 32)
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : None
...
```

show interfaces extensive (DS0 on Channelized T3 IQ)

```
user@host> show interfaces extensive ds-0/0/1:2:1
Physical interface: ds-0/0/1:2:1, Enabled, Physical link is Up
  Interface index: 176, SNMP ifIndex: 1563, Generation: 175
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: 640kbps,
  Loopback: None, FCS: 16, Parent: ct1-0/0/1:2(Index 175)
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
...
```


Copyright © 2018, Juniper Networks, Inc.

show interfaces (Discard)

Syntax
show interfaces dsc
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information Command introduced before Junos OS Release 7.4.

Description Display status information about the specified discard interface.

Options

dsc—Display standard information about the specified discard interface.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—This option is not relevant for the discard interface and always shows a value of 0.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) This option is not relevant for the discard interface and always shows a value of 0.

Required Privilege Level view

Related Documentation

- show interfaces (ATM) on page 1198
- *show interfaces routing*

List of Sample Output

show interfaces dsc on page 1324
show interfaces dsc brief on page 1324
show interfaces dsc detail on page 1324
show interfaces dsc extensive on page 1325

Output Fields Table 65 on page 1321 lists the output fields for the **show interfaces** (discard) command. Output fields are listed in the approximate order in which they appear.

Table 65: Discard show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		

Table 65: Discard show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
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
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 1110.	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 1110.	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 1110.	**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**

 Copyright © 2018, Juniper Networks, Inc.

Table 65: Discard show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes, 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	Input errors on the interface: • **Errors**—Sum of incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Giants**—Number of frames received that are larger than the giant threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • **Resource errors**—Sum of transmit drops.	**detail extensive**
Output errors	(Extensive only) Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **MTU errors**—Number of packets 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**

Copyright © 2018, Juniper Networks, Inc.

Table 65: 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 1110.	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
```


Copyright © 2018, Juniper Networks, Inc.

```
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 extensive

```
user@host> show interfaces dsc extensive
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
  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 (Fast Ethernet)

Syntax

show interfaces *interface-type*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display status information about the specified Fast Ethernet interface.

Options

interface-type—On M Series and T Series routers, the interface type is **fe-*fpc/pic/port***.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

List of Sample Output

show interfaces (Fast Ethernet) on page 1340
show interfaces brief (Fast Ethernet) on page 1340
show interfaces detail (Fast Ethernet) on page 1340
show interfaces extensive (Fast Ethernet) on page 1341

Output Fields

Table 66 on page 1326 lists the output fields for the **show interfaces** (Fast Ethernet) command. Output fields are listed in the approximate order in which they appear.

Table 66: show interfaces Fast Ethernet Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under “Common Output Fields Description” on page 1110.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none

Copyright © 2018, Juniper Networks, Inc.

Table 66: show interfaces Fast Ethernet Output Fields (continued)

Field Name	Field Description	Level of Output
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Link-mode	Type of link connection configured for the physical interface: **Full-duplex** or **Half-duplex**	**extensive**
Speed	Speed at which the interface is running.	All levels
Loopback	Loopback status: **Enabled** or **Disabled**. If loopback is enabled, type of loopback: **Local** or **Remote**.	All levels
Source filtering	Source filtering status: **Enabled** or **Disabled**.	All levels
LAN-PHY mode	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
WAN-PHY mode	10-Gigabit Ethernet interface operating in Wide Area Network Physical Layer Device (WAN PHY) mode. WAN PHY allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and other devices intended for SONET/SDH.	All levels
Unidirectional	Unidirectional link mode status for 10-Gigabit Ethernet interface: **Enabled** or **Disabled** for parent interface; **Rx-only** or **Tx-only** for child interfaces.	All levels
Flow control	Flow control status: **Enabled** or **Disabled**.	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: **Enabled** or **Disabled**.	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: • **Online**—Autonegotiation is manually configured as online. • **Offline**—Autonegotiation is manually configured as offline.	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under “Common Output Fields Description” on page 1110.	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 1110.	All levels
Link flags	Information about the link. Possible values are described in the “Links Flags” section under “Common Output Fields Description” on page 1110.	All levels
Wavelength	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels

Copyright © 2018, Juniper Networks, Inc.

Table 66: show interfaces Fast Ethernet Output Fields (continued)

Field Name	Field Description	Level of Output
Frequency	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels
CoS queues	Number of CoS queues configured.	**detail extensive** none
Schedulers	(GigabitEthernet intelligent queuing 2 (IQ2) interfaces only) Number of CoS schedulers configured.	**extensive**
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive** none
Hardware address	Hardware MAC address.	**detail extensive** none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour*:*minute*:*second*:*timezone*** **(*hour*:*minute*:*second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.	**detail extensive** none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface. Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes. For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 31 under the show interfaces command.	**detail extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 66: show interfaces Fast Ethernet Output Fields (continued)

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • **L3 incompletes**—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the **ignore-l3-incompletes** statement. • **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **FIFO errors**—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 66: show interfaces Fast Ethernet Output Fields (continued)

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Collisions**—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • **Aged packets**—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • **FIFO errors**—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**
Egress queues	Total number of egress queues supported on the specified interface. NOTE: In DPCs that are not of the enhanced type, such as DPC 40x 1GE R, DPCE 20x 1GE + 2x 10GE R, or DPCE 40x 1GE R, you might notice a discrepancy in the output of the **show interfaces** command because incoming packets might be counted in the Egress queues section of the output. This problem occurs on non-enhanced DPCs because the egress queue statistics are polled from IMQ (Inbound Message Queuing) block of the I-chip. The IMQ block does not differentiate between ingress and egress WAN traffic; as a result, the combined statistics are displayed in the egress queue counters on the Routing Engine. In a simple VPLS scenorio, if there is no MAC entry in DMAC table (by sending unidirectional traffic), traffic is flooded and the input traffic is accounted in IMQ. For bidirectional traffic (MAC entry in DMAC table), if the outgoing interface is on the same I-chip then both ingress and egress statistics are counted in a combined way. If the outgoing interface is on a different I-chip or FPC, then only egress statistics are accounted in IMQ. This behavior is expected with non-enhanced DPCs	**detail extensive**
Queue counters (Egress)	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 66: show interfaces Fast Ethernet Output Fields (continued)

Field Name	Field Description	Level of Output
Ingress queues	Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.	**extensive**
Queue counters (Ingress)	CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**extensive**
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the routing device configuration, an alarm can ring the red or yellow alarm bell on the routing device, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value **None** or **Link**. • **None**—There are no active defects or alarms. • **Link**—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	**detail extensive** none
OTN FEC statistics	The forward error correction (FEC) counters provide the following statistics:. • **Corrected Errors**—The count of corrected errors in the last second. • **Corrected Error Ratio**—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits.	
PCS statistics	(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device. • **Bit errors**—The number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode. • **Errored blocks**—The number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode.	**detail extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 66: show interfaces Fast Ethernet Output Fields (continued)

Field Name	Field Description	Level of Output
MAC statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC subsystem, including the following: • **Total octets** and **total packets**—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see Table 31 under the show interfaces command. • **Unicast packets, Broadcast packets,** and **Multicast packets**—Number of unicast, broadcast, and multicast packets. • **CRC/Align errors**—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • **FIFO error**—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • **MAC control frames**—Number of MAC control frames. • **MAC pause frames**—Number of MAC control frames with **pause** operational code. • **Oversized frames**—Number of frames that exceed 1518 octets. • **Jabber frames**—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • **Fragment frames**—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • **VLAN tagged frames**—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. • **Code violations**—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error."	**extensive**
OTN Received Overhead Bytes	APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08	**extensive**
OTN Transmitted Overhead Bytes	APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08	**extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 66: show interfaces Fast Ethernet Output Fields (continued)

Field Name	Field Description	Level of Output
Filter statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected. • **Input packet count**—Number of packets received from the MAC hardware that the filter processed. • **Input packet rejects**—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • **Input DA rejects**—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the routing device from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local routing device (which the routing device is rejecting). • **Input SA rejects**—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • **Output packet count**—Number of packets that the filter has given to the MAC hardware. • **Output packet pad count**—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • **Output packet error count**—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment. • **CAM destination filters, CAM source filters**—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields should be 0.	**extensive**
PMA PHY	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information: • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. Any state other than **OK** indicates a problem. Subfields are: • **PHY Lock**—Phase-locked loop • **PHY Light**—Loss of optical signal	**extensive**

Table 66: show interfaces Fast Ethernet Output Fields (continued)

Field Name	Field Description	Level of Output
WIS section	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information: • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. Any state other than **OK** indicates a problem. Subfields are: • **BIP-B1**—Bit interleaved parity for SONET section overhead • **SEF**—Severely errored framing • **LOL**—Loss of light • **LOF**—Loss of frame • **ES-S**—Errored seconds (section) • **SES-S**—Severely errored seconds (section) • **SEFS-S**—Severely errored framing seconds (section)	**extensive**
WIS line	(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **BIP-B2**—Bit interleaved parity for SONET line overhead • **REI-L**—Remote error indication (near-end line) • **RDI-L**—Remote defect indication (near-end line) • **AIS-L**—Alarm indication signal (near-end line) • **BERR-SF**—Bit error rate fault (signal failure) • **BERR-SD**—Bit error rate defect (signal degradation) • **ES-L**—Errored seconds (near-end line) • **SES-L**—Severely errored seconds (near-end line) • **UAS-L**—Unavailable seconds (near-end line) • **ES-LFE**—Errored seconds (far-end line) • **SES-LFE**—Severely errored seconds (far-end line) • **UAS-LFE**—Unavailable seconds (far-end line)	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 66: show interfaces Fast Ethernet Output Fields (continued)

Field Name	Field Description	Level of Output
WIS path	(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. Any state other than **OK** indicates a problem. Subfields are: • **BIP-B3**—Bit interleaved parity for SONET section overhead • **REI-P**—Remote error indication • **LOP-P**—Loss of pointer (path) • **AIS-P**—Path alarm indication signal • **RDI-P**—Path remote defect indication • **UNEQ-P**—Path unequipped • **PLM-P**—Path payload (signal) label mismatch • **ES-P**—Errored seconds (near-end STS path) • **SES-P**—Severely errored seconds (near-end STS path) • **UAS-P**—Unavailable seconds (near-end STS path) • **SES-PFE**—Severely errored seconds (far-end STS path) • **UAS-PFE**—Unavailable seconds (far-end STS path)	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 66: show interfaces Fast Ethernet Output Fields (continued)

Field Name	Field Description	Level of Output

Autonegotiation information	Information about link autonegotiation.
	• **Negotiation status:**
	  ◦ **Incomplete**—Ethernet interface has the speed or link mode configured.
	  ◦ **No autonegotiation**—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
	  ◦ **Complete**—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• **Link partner status**—**OK** when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• **Link partner:**
	  ◦ **Link mode**—Depending on the capability of the attached Ethernet device, either **Full-duplex** or **Half-duplex**.
	  ◦ **Flow control**—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is **None**. For Gigabit Ethernet interfaces, types are **Symmetric** (link partner supports **PAUSE** on receive and transmit), **Asymmetric** (link partner supports **PAUSE** on transmit), and **Symmetric/Asymmetric** (link partner supports both **PAUSE** on receive and transmit or only **PAUSE** receive).
	  ◦ **Remote fault**—Remote fault information from the link partner—**Failure** indicates a receive link error. **OK** indicates that the link partner is receiving. **Negotiation error** indicates a negotiation error. **Offline** indicates that the link partner is going offline.
	• **Local resolution**—Information from the link partner:
	  ◦ **Flow control**—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are **Symmetric** (link partner supports **PAUSE** on receive and transmit), **Asymmetric** (link partner supports **PAUSE** on transmit), and **Symmetric/Asymmetric** (link partner supports both **PAUSE** on receive and transmit or only **PAUSE** receive).
	  ◦ **Remote fault**—Remote fault information. **Link OK** (no error detected on receive), **Offline** (local interface is offline), and **Link Failure** (link error detected on receive).
		extensive

Received path trace, Transmitted path trace	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other routing device manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the routing device at the other end of the fiber. The transmitted path trace value is the message that this routing device transmits.	**extensive**
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number.	**extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 66: show interfaces Fast Ethernet Output Fields (continued)

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface. • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP interface index number for the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under "Common Output Fields Description" on page 1110.	All levels

Table 66: show interfaces Fast Ethernet Output Fields (continued)

Field Name	Field Description	Level of Output
VLAN-Tag	Rewrite profile applied to incoming or outgoing frames on the outer (**Out**) VLAN tag or for both the outer and inner (**In**) VLAN tags. • **push**—An outer VLAN tag is pushed in front of the existing VLAN tag. • **pop**—The outer VLAN tag of the incoming frame is removed. • **swap**—The outer VLAN tag of the incoming frame is overwritten with the user specified VLAN tag information. • **push**—An outer VLAN tag is pushed in front of the existing VLAN tag. • **push-push**—Two VLAN tags are pushed in from the incoming frame. • **swap-push**—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • **swap-swap**—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user specified VLAN tag value. • **pop-swap**—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • **pop-pop**—Both the outer and inner VLAN tags of the incoming frame are removed.	**brief detail extensive** none
Demux:	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: • Source Family Inet • Destination Family Inet	**detail extensive** none
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under “Common Output Fields Description” on page 1110.	**detail extensive** none
MTU	Maximum transmission unit size on the logical interface.	**detail extensive** none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	**detail extensive** none
Traffic statistics	Number and rate of bytes and packets received and transmitted on the specified interface set. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface set • **Input packets, Output packets**—Number of packets received and transmitted on the interface set.	**detail extensive**
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	**extensive**
Local statistics	Number and rate of bytes and packets destined to the routing device.	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 66: show interfaces Fast 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 1110.	**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 1110.	**detail extensive** none
protocol-family	Protocol family configured on the logical interface. If the protocol is **inet**, the IP address of the interface is also displayed.	**brief**
Flags	Information about address flag (possible values are described in the "Addresses Flags" section under "Common Output Fields Description" on page 1110.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interlace.	**detail extensive** none

Copyright © 2018, Juniper Networks, Inc.

Table 66: show interfaces Fast Ethernet Output Fields (continued)

Field Name	Field Description	Level of Output
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces (Fast Ethernet)

```
user@host> show interfaces fe-0/0/0
Physical interface: fe-0/0/0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 22
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:00:5e:00:53:38, Hardware address: 00:00:5e:00:53:38
  Last flapped   : 2006-01-20 14:50:58 PST (2w4d 00:44 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 203.0.113/24, Local: 203.0.113.1, Broadcast: 203.0.113.255
```

show interfaces brief (Fast Ethernet)

```
user@host> show interfaces fe-0/0/0 brief
Physical interface: fe-0/0/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Logical interface fe-0/0/0.0
    Flags: SNMP-Traps Encapsulation: ENET2
    inet  203.0.113.1/24
```

show interfaces detail (Fast Ethernet)

```
user@host> show interfaces fe-0/0/0 detail
Physical interface: fe-0/0/0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 22, Generation: 5391
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:38, Hardware address: 00:00:5e:00:53:3f:38
  Last flapped   : 2006-01-20 14:50:58 PST (2w4d 00:45 ago)
  Statistics last cleared: Never
  Traffic statistics:
```


Copyright © 2018, Juniper Networks, Inc.

```
     Input  bytes  :                  0                    0 bps
     Output bytes  :                 42                    0 bps
     Input  packets:                  0                    0 pps
     Output packets:                  1                    0 pps
    Active alarms  : None
    Active defects : None
    Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198) (Generation 67)
      Flags: SNMP-Traps Encapsulation: ENET2
      Protocol inet, MTU: 1500, Generation: 105, Route table: 0
        Flags: Is-Primary, Mac-Validate-Strict
        Mac-Validate Failures: Packets: 0, Bytes: 0
        Addresses, Flags: Is-Preferred Is-Primary
        Destination: 203.0.113/24, Local: 203.0.113.1, Broadcast: 203.0.113.255,
          Generation: 136
```

show interfaces extensive (Fast Ethernet)

```
user@host> show interfaces fe-0/0/0 extensive
Physical interface: fe-0/0/0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 22, Generation: 5391
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed:
  100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:38, Hardware address: 00:00:5e:00:53:38
  Last flapped   : 2006-01-20 14:50:58 PST (2w4d 00:46 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                  0                    0 bps
   Output bytes  :                 42                    0 bps
   Input  packets:                  0                    0 pps
   Output packets:                  1                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 3, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Active alarms  : None
  Active defects : None
  MAC statistics:                      Receive         Transmit
    Total octets                             0               64
    Total packets                            0                1
    Unicast packets                          0                0
    Broadcast packets                        0                1
    Multicast packets                        0                0
    CRC/Align errors                         0                0
    FIFO errors                              0                0
    MAC control frames                       0                0
    MAC pause frames                         0                0
    Oversized frames                         0
    Jabber frames                            0
    Fragment frames                          0
    VLAN tagged frames                       0
    Code violations                          0
Filter statistics:
```

```
    Input packet count                      0
    Input packet rejects                    0
    Input DA rejects                        0
    Input SA rejects                        0
    Output packet count                                      1
    Output packet pad count                                  0
    Output packet error count                                0
    CAM destination filters: 1, CAM source filters: 0
  Autonegotiation information:
    Negotiation status: Complete
    Link partner:
        Link partner: Full-duplex, Flow control: None, Remote fault: Ok
    Local resolution:
Packet Forwarding Engine configuration:
    Destination slot: 0
CoS information:
                 Bandwidth           Buffer Priority   Limit
                              %           bps   %         usec
    0 best-effort            95    950000000  95             0      low    none
    3 network-control         5     50000000   5             0      low    none
  Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198) (Generation 67)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Generation: 105, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 203.0.113/24, Local: 203.0.113.1, Broadcast: 203.0.113.255,

        Generation: 136
```

 Copyright © 2018, Juniper Networks, Inc.

show interfaces

List of Syntax

Syntax (Gigabit Ethernet) on page 1343
Syntax (10 Gigabit Ethernet) on page 1343
Syntax (SRX Series Devices) on page 1343

Syntax (Gigabit Ethernet)

show interfaces ge-*fpc/pic/port*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Syntax (10 Gigabit Ethernet)

show interfaces xe-*fpc/pic/port*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Syntax (SRX Series Devices)

show interfaces (
<*interface-name*>
<brief | detail | extensive | terse>
<controller *interface-name*>|
<descriptions *interface-name*>|
<destination-class (all | *destination-class-name logical-interface-name*)>|
<diagnostics optics *interface-name*>|
<far-end-interval *interface-fpc/pic/port*>|
<filters *interface-name*>|
<flow-statistics *interface-name*>|
<interval *interface-name*>|
<load-balancing (detail | *interface-name*)>|
<mac-database mac-address *mac-address*>|
<mc-ae id *identifier* unit *number* revertive-info>|
<media *interface-name*>|
<policers *interface-name*>|
<queue both-ingress-egress egress forwarding-class *forwarding-class* ingress l2-statistics>|
<redundancy (detail | *interface-name*)>|
<routing brief detail summary *interface-name*>|
<routing-instance (all | *instance-name*)>|
<snmp-index *snmp-index*>|
<source-class (all | *destination-class-name logical-interface-name*)>|
<statistics *interface-name*>|
<switch-port *switch-port number*>|
<transport pm (all | optics | otn) (all | current | currentday | interval | previousday) (all | *interface-name*)>|
<zone *interface-name*>
)

Release Information

Command introduced before Junos OS Release 7.4 for Gigabit interfaces.
Command introduced in Junos OS Release 8.0 for 10 Gigabit interfaces.
Command modified in Junos OS Release 9.5 for SRX Series devices.

Copyright © 2018, Juniper Networks, Inc.

Command introduced in Junos OS Release 18.1 for Gigabit interfaces.

Description Display status information about the specified Gigabit Ethernet interface.

(M320, M120, MX Series, and T Series routers only) Display status information about the specified 10-Gigabit Ethernet interface.

Display the IPv6 interface traffic statistics about the specified Gigabit Ethernet interface for MX series routers. The input and output bytes (bps) and packets (pps) rates are not displayed for IFD and local traffic.

Display status information and statistics about interfaces on SRX Series appliance running Junos OS.

NOTE: On SRX Series appliances, on configuring identical IPs on a single interface, you will not see a warning message; instead, you will see a syslog message.

Options For Gigabit interfaces:

ge-*fpc/pic/port*—Display standard information about the specified Gigabit Ethernet interface.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

For 10 Gigabit interfaces:

xe-*fpc/pic/port*—Display standard information about the specified 10-Gigabit Ethernet interface.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

For SRX interfaces:

Copyright © 2018, Juniper Networks, Inc.

- **interface-name**—(Optional) Display standard information about the specified interface. Following is a list of typical interface names. Replace pim with the PIM slot and port with the port number.
  - **at-** ***pim/0/port***—ATM-over-ADSL or ATM-over-SHDSL interface.
  - **ce1-*pim/0/ port***—Channelized E1 interface.
  - **cl-0/0/8**—3G wireless modem interface for SRX320 devices.
  - **ct1-*pim/0/port***—Channelized T1 interface.
  - **dl0**—Dialer Interface for initiating ISDN and USB modem connections.
  - **e1-*pim/0/port***—E1 interface.
  - **e3-*pim/0/port***—E3 interface.
  - **fe-*pim/0/port***—Fast Ethernet interface.
  - **ge-*pim/0/port***—Gigabit Ethernet interface.
  - **se-*pim/0/port***—Serial interface.
  - **t1-*pim/0/port***—T1 (also called DS1) interface.
  - **t3-*pim/0/port***—T3 (also called DS3) interface.
  - **wx-*slot*/0/0**—WAN acceleration interface, for the WXC Integrated Services Module (ISM 200).
  - **interface-name**—(Optional) Display standard information about the specified interface. Following is a list of typical interface names. Replace pim with the PIM slot and port with the port number.
    - **at-** ***pim/0/port***—ATM-over-ADSL or ATM-over-SHDSL interface.
    - **ce1-*pim/0/ port***—Channelized E1 interface.
    - **cl-0/0/8**—3G wireless modem interface for SRX320 devices.
    - **ct1-*pim/0/port***—Channelized T1 interface.
    - **dl0**—Dialer Interface for initiating ISDN and USB modem connections.
    - **e1-*pim/0/port***—E1 interface.
    - **e3-*pim/0/port***—E3 interface.
    - **fe-*pim/0/port***—Fast Ethernet interface.
    - **ge-*pim/0/port***—Gigabit Ethernet interface.
    - **se-*pim/0/port***—Serial interface.
    - **t1-*pim/0/port***—T1 (also called DS1) interface.

- **t3-*pim*/0/*port***—T3 (also called DS3) interface.
- **wx-*slot*/0/0**—WAN acceleration interface, for the WXC Integrated Services Module (ISM 200).

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

- *Understanding Layer 2 Interfaces on Security Devices*
- *Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*
- *Verifying and Managing Configurations for Dynamic VLANs Based on Access-Line Identifiers*

List of Sample Output

Copyright © 2018, Juniper Networks, Inc.

Output Fields Table 56 on page 1123 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 57 on page 1152.

Table 67: 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 1110.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none

Table 67: 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**
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Loopback	Loopback status: **Enabled** or **Disabled**. If loopback is enabled, type of loopback: **Local** or **Remote**.	All levels
Source filtering	Source filtering status: **Enabled** or **Disabled**.	All levels
LAN-PHY mode	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
WAN-PHY mode	10-Gigabit Ethernet interface operating in Wide Area Network Physical Layer Device (WAN PHY) mode. WAN PHY allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and other devices intended for SONET/SDH.	All levels
Unidirectional	Unidirectional link mode status for 10-Gigabit Ethernet interface: **Enabled** or **Disabled** for parent interface; **Rx-only** or **Tx-only** for child interfaces.	All levels
Flow control	Flow control status: **Enabled** or **Disabled**.	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: **Enabled** or **Disabled**.	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: • **Online**—Autonegotiation is manually configured as online. • **Offline**—Autonegotiation is manually configured as offline.	All levels
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under "Common Output Fields Description" on page 1110.	All levels
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under "Common Output Fields Description" on page 1110.	All levels
Link flags	Information about the link. Possible values are described in the "Links Flags" section under "Common Output Fields Description" on page 1110.	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

 Copyright © 2018, Juniper Networks, Inc.

Table 67: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
Schedulers	(Gigabit Ethernet intelligent queuing 2 [IQ2] interfaces only) Number of CoS schedulers configured.	**extensive**
Hold-times	Current interface hold-time up and hold-time down, in milliseconds (ms).	**detail extensive**
Current address	Configured MAC address.	**detail extensive** none
Hardware address	Hardware MAC address.	**detail extensive** none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour:minute:second:timezone*** **(*hour:minute:second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).**	**detail extensive** none
Input Rate	Input rate in bits per second (bps) and packets per second (pps). The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None
Output Rate	Output rate in bps and pps. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Egress account overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for egress traffic.	**detail extensive**
Ingress account overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for ingress traffic.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Output bytes**—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface. Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes. For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 31 under the show interfaces command.	**detail extensive**

Table 67: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • **L3 incompletes**—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the **ignore-l3-incompletes** statement. • **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **FIFO errors**—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **Resource errors**—Sum of transmit drops.	**extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 67: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the **Drops** field does not always use the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs. • **Collisions**—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number must always be 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 must 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. NOTE: In DPCs that are not of the enhanced type, such as DPC 40x 1GE R, DPCE 20x 1GE + 2x 10GE R, or DPCE 40x 1GE R, you might notice a discrepancy in the output of the **show interfaces** command because incoming packets might be counted in the Egress queues section of the output. This problem occurs on non-enhanced DPCs because the egress queue statistics are polled from IMQ (Inbound Message Queuing) block of the I-chip. The IMQ block does not differentiate between ingress and egress WAN traffic; as a result, the combined statistics are displayed in the egress queue counters on the Routing Engine. In a simple VPLS scenario, if there is no MAC entry in DMAC table (by sending unidirectional traffic), traffic is flooded and the input traffic is accounted in IMQ. For bidirectional traffic (MAC entry in DMAC table), if the outgoing interface is on the same I-chip then both ingress and egress statistics are counted in a combined way. If the outgoing interface is on a different I-chip or FPC, then only egress statistics are accounted in IMQ. This behavior is expected with non-enhanced DPCs	**detail extensive**

Table 67: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
Queue counters (Egress)	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism. NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the **Dropped packets** field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.	**detail extensive**
Ingress queues	Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.	**extensive**
Queue counters (Ingress)	CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**extensive**
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value **None** or **Link**. • **None**—There are no active defects or alarms. • **Link**—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	**detail extensive** none
Interface transmit statistics	(On MX Series devices) Status of the **interface-transmit-statistics** configuration: Enabled or Disabled. • **Enabled**—When the **interface-transmit-statistics** statement is included in the configuration. If this is configured, the interface statistics show the actual transmitted load on the interface. • **Disabled**—When the **interface-transmit-statistics** statement is not included in the configuration. If this is not configured, the interface statistics show the offered load on the interface.	**detail extensive**
OTN FEC statistics	The forward error correction (FEC) counters provide the following statistics: • **Corrected Errors**—Count of corrected errors in the last second. • **Corrected Error Ratio**—Corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits.	**detail extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 67: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
PCS statistics	(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device. • **Bit errors**—Number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode. • **Errored blocks**—Number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode.	**detail extensive**

Link Degrade	Shows the link degrade status of the physical link and the estimated bit error rates (BERs). This field is available only for the PICs supporting the physical link monitoring feature.
	• **Link Monitoring**—Indicates if physical link degrade monitoring is enabled on the interface.
	  ◦ **Enable**—Indicates that link degrade monitoring has been enabled (using the **link-degrade-monitor** statement) on the interface.
	  ◦ **Disable**—Indicates that link degrade monitoring has not been enabled on the interface. If link degrade monitoring has not been enabled, the output does not show any related information, such as BER values and thresholds.
	• **Link Degrade Set Threshold**—The BER threshold value at which the link is considered degraded and a corrective action is triggered.
	• **Link Degrade Clear Threshold**—The BER threshold value at which the degraded link is considered recovered and the corrective action applied to the interface is reverted.
	• **Estimated BER**—The estimated bit error rate.
	• **Link-degrade event**—Shows link degrade event information.
	  ◦ **Seconds**—Time (in seconds) elapsed after a link degrade event occurred.
	  ◦ **Count**—The number of link degrade events recorded.
	  ◦ **State**—Shows the link degrade status (example: **Defect Active**).
		detail extensive

Table 67: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output

MAC statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC subsystem, including the following:
	• **Total octets** and **total packets**—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see Table 31 under the show interfaces command.
	• **Unicast packets, Broadcast packets,** and **Multicast packets**—Number of unicast, broadcast, and multicast packets.
	• **CRC/Align errors**—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).
	• **FIFO error**—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning.
	• **MAC control frames**—Number of MAC control frames.
	• **MAC pause frames**—Number of MAC control frames with **pause** operational code.
	• **Oversized frames**—There are two possible conditions regarding the number of oversized frames:
	  ◦ Packet length exceeds 1518 octets, or
	  ◦ Packet length exceeds MRU
	• **Jabber frames**—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.
	• **Fragment frames**—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.
	• **VLAN tagged frames**—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not.
	NOTE: The 20-port Gigabit Ethernet MIC (MIC-3D-20GE-SFP) does not have hardware counters for VLAN frames. Therefore, the **VLAN tagged frames** field displays 0 when the **show interfaces** command is executed on a 20-port Gigabit Ethernet MIC. In other words, the number of VLAN tagged frames cannot be determined for the 20-port Gigabit Ethernet MIC.
	• **Code violations**—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”
		extensive

OTN Received Overhead Bytes	APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08	**extensive**
OTN Transmitted Overhead Bytes	APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08	**extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 67: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
Filter statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet may enter the system or be rejected. • **Input packet count**—Number of packets received from the MAC hardware that the filter processed. • **Input packet rejects**—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • **Input DA rejects**—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the router from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local router (which the router is rejecting). • **Input SA rejects**—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field must 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 must 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 must be 0.	**extensive**
PMA PHY	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information: • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. Any state other than **OK** indicates a problem. Subfields are: • **PHY Lock**—Phase-locked loop • **PHY Light**—Loss of optical signal	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 67: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
WIS section	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information: • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. Any state other than **OK** indicates a problem. Subfields are: • **BIP-B1**—Bit interleaved parity for SONET section overhead • **SEF**—Severely errored framing • **LOL**—Loss of light • **LOF**—Loss of frame • **ES-S**—Errored seconds (section) • **SES-S**—Severely errored seconds (section) • **SEFS-S**—Severely errored framing seconds (section)	**extensive**
WIS line	(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information: • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. Any state other than **OK** indicates a problem. Subfields are: • **BIP-B2**—Bit interleaved parity for SONET line overhead • **REI-L**—Remote error indication (near-end line) • **RDI-L**—Remote defect indication (near-end line) • **AIS-L**—Alarm indication signal (near-end line) • **BERR-SF**—Bit error rate fault (signal failure) • **BERR-SD**—Bit error rate defect (signal degradation) • **ES-L**—Errored seconds (near-end line) • **SES-L**—Severely errored seconds (near-end line) • **UAS-L**—Unavailable seconds (near-end line) • **ES-LFE**—Errored seconds (far-end line) • **SES-LFE**—Severely errored seconds (far-end line) • **UAS-LFE**—Unavailable seconds (far-end line)	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 67: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
WIS path	(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information: • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. Any state other than **OK** indicates a problem. Subfields are: • **BIP-B3**—Bit interleaved parity for SONET section overhead • **REI-P**—Remote error indication • **LOP-P**—Loss of pointer (path) • **AIS-P**—Path alarm indication signal • **RDI-P**—Path remote defect indication • **UNEQ-P**—Path unequipped • **PLM-P**—Path payload (signal) label mismatch • **ES-P**—Errored seconds (near-end STS path) • **SES-P**—Severely errored seconds (near-end STS path) • **UAS-P**—Unavailable seconds (near-end STS path) • **SES-PFE**—Severely errored seconds (far-end STS path) • **UAS-PFE**—Unavailable seconds (far-end STS path)	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 67: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output

Autonegotiation information	Information about link autonegotiation.
	• **Negotiation status:**
	  ◦ **Incomplete**—Ethernet interface has the speed or link mode configured.
	  ◦ **No autonegotiation**—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
	  ◦ **Complete**—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• **Link partner status**—**OK** when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• **Link partner**—Information from the remote Ethernet device:
	  ◦ **Link mode**—Depending on the capability of the link partner, either **Full-duplex** or **Half-duplex**.
	  ◦ **Flow control**—Types of flow control supported by the link partner. For Gigabit Ethernet interfaces, types are **Symmetric** (link partner supports **PAUSE** on receive and transmit), **Asymmetric** (link partner supports **PAUSE** on transmit), **Symmetric/Asymmetric** (link partner supports **PAUSE** on receive and transmit or only **PAUSE** on transmit), and **None** (link partner does not support flow control).
	  ◦ **Remote fault**—Remote fault information from the link partner—**Failure** indicates a receive link error. **OK** indicates that the link partner is receiving. **Negotiation error** indicates a negotiation error. **Offline** indicates that the link partner is going offline.
	• **Local resolution**—Information from the local Ethernet device:
	  ◦ **Flow control**—Types of flow control supported by the local device. For Gigabit Ethernet interfaces, advertised capabilities are **Symmetric/Asymmetric** (local device supports **PAUSE** on receive and transmit or only **PAUSE** on receive) and **None** (local device does not support flow control). Depending on the result of the negotiation with the link partner, local resolution flow control type will display **Symmetric** (local device supports **PAUSE** on receive and transmit), **Asymmetric** (local device supports **PAUSE** on receive), and **None** (local device does not support flow control).
	  ◦ **Remote fault**—Remote fault information. **Link OK** (no error detected on receive), **Offline** (local interface is offline), and **Link Failure** (link error detected on receive).
		extensive

Received path trace, Transmitted path trace	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the router at the other end of the fiber. The transmitted path trace value is the message that this router transmits.	**extensive**
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number.	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 67: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface. • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP interface index number for the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under "Common Output Fields Description" on page 1110.	All levels

Table 67: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
VLAN-Tag	Rewrite profile applied to incoming or outgoing frames on the outer (**Out**) VLAN tag or for both the outer and inner (**In**) VLAN tags. • **push**—An outer VLAN tag is pushed in front of the existing VLAN tag. • **pop**—The outer VLAN tag of the incoming frame is removed. • **swap**—The outer VLAN tag of the incoming frame is overwritten with the user-specified VLAN tag information. • **push**—An outer VLAN tag is pushed in front of the existing VLAN tag. • **push-push**—Two VLAN tags are pushed in from the incoming frame. • **swap-push**—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • **swap-swap**—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user-specified VLAN tag value. • **pop-swap**—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • **pop-pop**—Both the outer and inner VLAN tags of the incoming frame are removed.	**brief detail extensive** none
Demux	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: • Source Family Inet • Destination Family Inet	**detail extensive** none
Encapsulation	Encapsulation on the logical interface.	All levels
ACI VLAN	Information displayed for agent circuit identifier (ACI) interface set configured with the **agent-circuit-id** autoconfiguration stanza. **Dynamic Profile**—Name of the dynamic profile that defines the ACI interface set. If configured, the ACI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ACI information. NOTE: The ACI VLAN field is replaced with the Line Identity field when an ALI interface set is configured with the **line-identity** autoconfiguration stanza.	**brief detail extensive** none
Line Identity	Information displayed for access-line-identifier (ALI) interface sets configured with the **line-identity** autoconfiguration stanza. • **Dynamic Profile**—Name of the dynamic profile that defines the ALI interface set. • Trusted option used to create the ALI interface set: **Circuit-id**, **Remote-id**, or **Accept-no-ids**. More than one option can be configured. If configured, the ALI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ALI information. NOTE: The Line Identity field is replaced with the ACI VLAN field when an ACI interface set is configured with the **agent-circuit-id** autoconfiguration stanza.	**detail**

 Copyright © 2018, Juniper Networks, Inc.

Table 67: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under “Common Output Fields Description” on page 1110.	**detail extensive** none
MTU	Maximum transmission unit size on the logical interface.	**detail extensive** none
Neighbor Discovery Protocol (NDP)Queue Statistics	NDP statistics for protocol **inet6** under logical interface statistics. • **Max nh cache**—Maximum interface neighbor discovery nexthop cache size. • **New hold nh limit**—Maximum number of new unresolved nexthops. • **Curr nh cnt**—Current number of resolved nexthops in the NDP queue. • **Curr new hold cnt**—Current number of unresolved nexthops in the NDP queue. • **NH drop cnt**—Number of NDP requests not serviced.	All levels
Dynamic Profile	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	Name of the service name table for the interface configured with a PPPoE family.	**detail extensive** none
Max Sessions	Maximum number of PPPoE logical interfaces that can be activated on the underlying interface.	**detail extensive** none
Duplicate Protection	State of PPPoE duplicate protection: **On** or **Off**. When duplicate protection is configured for the underlying interface, a dynamic PPPoE logical interface cannot be activated when an existing active logical interface is present for the same PPPoE client.	**detail extensive** none
Direct Connect	State of the configuration to ignore DSL Forum VSAs: **On** or **Off**. When configured, the router ignores any of these VSAs received from a directly connected CPE device on the interface.	**detail extensive** none
AC Name	Name of the access concentrator.	**detail extensive** none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	**detail extensive** none
Traffic statistics	Number and rate of bytes and packets received and transmitted on the specified interface set. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Input packets, Output packets**—Number of packets received and transmitted on the interface set.	**detail extensive**
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	**extensive**
Local statistics	Number and rate of bytes and packets destined to the router.	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 67: 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 1110.	**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**
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 1110.	**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 1110.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none

 Copyright © 2018, Juniper Networks, Inc.

Table 67: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
Broadcast	Broadcast address of the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

The following table describes the output fields for the **show interfaces** (10–Gigabit Ethernet) command.

Field Name	Field Description	Level of Output
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 1110.	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

Copyright © 2018, Juniper Networks, Inc.

Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: • **Online**—Autonegotiation is manually configured as online. • **Offline**—Autonegotiation is manually configured as offline.	All levels
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under "Common Output Fields Description" on page 1110.	All levels
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under "Common Output Fields Description" on page 1110.	All levels
Link flags	Information about the link. Possible values are described in the "Links Flags" section under "Common Output Fields Description" on page 1110.	All levels
Wavelength	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels
Frequency	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels
CoS queues	Number of CoS queues configured.	**detail extensive** none
Schedulers	(Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces only) Number of CoS schedulers configured.	**extensive**
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive** none
Hardware address	Hardware MAC address.	**detail extensive** none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour:minute:second:timezone*** **(*hour:minute:second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).**	**detail extensive** none
Input Rate	Input rate in bits per second (bps) and packets per second (pps). The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None specified
Output Rate	Output rate in bps and pps. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Egress account overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for egress traffic.	**detail extensive**
Ingress account overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for ingress traffic.	**detail extensive**

 Copyright © 2018, Juniper Networks, Inc.

Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Output bytes**—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • **L3 incompletes**—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the **ignore-l3-incompletes** statement. • **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **FIFO errors**—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **Resource errors**—Sum of transmit drops.	**extensive**

Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Collisions**—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • **Aged packets**—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • **FIFO errors**—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**
Egress queues	Total number of egress queues supported on the specified interface. NOTE: In DPCs that are not of the enhanced type, such as DPC 40x 1GE R, DPCE 20x 1GE + 2x 10GE R, or DPCE 40x 1GE R, you might notice a discrepancy in the output of the **show interfaces** command because incoming packets might be counted in the Egress queues section of the output. This problem occurs on non-enhanced DPCs because the egress queue statistics are polled from IMQ (Inbound Message Queuing) block of the I-chip. The IMQ block does not differentiate between ingress and egress WAN traffic; as a result, the combined statistics are displayed in the egress queue counters on the Routing Engine. In a simple VPLS scenorio, if there is no MAC entry in DMAC table (by sending unidirectional traffic), traffic is flooded and the input traffic is accounted in IMQ. For bidirectional traffic (MAC entry in DMAC table), if the outgoing interface is on the same I-chip then both ingress and egress statistics are counted in a combined way. If the outgoing interface is on a different I-chip or FPC, then only egress statistics are accounted in IMQ. This behavior is expected with non-enhanced DPCs	**detail extensive**
Queue counters (Egress)	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**
Ingress queues	Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Queue counters (Ingress)	CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**extensive**
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the routing device configuration, an alarm can ring the red or yellow alarm bell on the routing device, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value **None** or **Link**. • **None**—There are no active defects or alarms. • **Link**—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	**detail extensive** none
OTN alarms	Active OTN alarms identified on the interface.	**detail extensive**
OTN defects	OTN defects received on the interface.	**detail extensive**
OTN FEC Mode	The FECmode configured on the interface. • **efec**—Enhanced forward error correction (EFEC) is configured to defect and correct bit errors. • **gfec**—G.709 Forward error correction (GFEC) mode is configured to detect and correct bit errors. • **none**—FEC mode is not configured.	**detail extensive**
OTN Rate	OTN mode. • **fixed-stuff-bytes**—Fixed stuff bytes 11.0957 Gbps. • **no-fixed-stuff-bytes**—No fixed stuff bytes 11.0491 Gbps. • **pass-through**—Enable OTN passthrough mode. • **no-pass-through**—Do not enable OTN passthrough mode.	**detail extensive**
OTN Line Loopback	Status of the line loopback, if configured for the DWDM OTN PIC. Its value can be: **enabled** or **disabled**.	**detail extensive**
OTN FEC statistics	The forward error correction (FEC) counters for the DWDM OTN PIC. • **Corrected Errors**—The count of corrected errors in the last second. • **Corrected Error Ratio**—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits.	**detail extensive**
OTN FEC alarms	OTN FEC excessive or degraded error alarms triggered on the interface. • **FEC Degrade**—OTU FEC Degrade defect. • **FEC Excessive**—OTU FEC Excessive Error defect.	**detail extensive**

OTN OC	OTN OC defects triggered on the interface. • **LOS**—OC Loss of Signal defect. • **LOF**—OC Loss of Frame defect. • **LOM**—OC Loss of Multiframe defect. • **Wavelength Lock**—OC Wavelength Lock defect.	**detail extensive**
OTN OTU	OTN OTU defects detected on the interface • **AIS**—OTN AIS alarm. • **BDI**—OTN OTU BDI alarm. • **IAE**—OTN OTU IAE alarm. • **TTIM**—OTN OTU TTIM alarm. • **SF**—OTN ODU bit error rate fault alarm. • **SD**—OTN ODU bit error rate defect alarm. • **TCA-ES**—OTN ODU ES threshold alarm. • **TCA-SES**—OTN ODU SES threshold alarm. • **TCA-UAS**—OTN ODU UAS threshold alarm. • **TCA-BBE**—OTN ODU BBE threshold alarm. • **BIP**—OTN ODU BIP threshold alarm. • **BBE**—OTN OTU BBE threshold alarm. • **ES**—OTN OTU ES threshold alarm. • **SES**—OTN OTU SES threshold alarm. • **UAS**—OTN OTU UAS threshold alarm.	**detail extensive**
Received DAPI	Destination Access Port Interface (DAPI) from which the packets were received.	**detail extensive**
Received SAPI	Source Access Port Interface (SAPI) from which the packets were received.	**detail extensive**
Transmitted DAPI	Destination Access Port Interface (DAPI) to which the packets were transmitted.	**detail extensive**
Transmitted SAPI	Source Access Port Interface (SAPI) to which the packets were transmitted.	**detail extensive**
PCS statistics	(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device. • **Bit errors**—The number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode. • **Errored blocks**—The number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode.	**detail extensive**

 Copyright © 2018, Juniper Networks, Inc.

MAC statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC subsystem, including the following: • **Total octets** and **total packets**—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. • **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

Filter statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected. • **Input packet count**—Number of packets received from the MAC hardware that the filter processed. • **Input packet rejects**—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • **Input DA rejects**—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the routing device from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local routing device (which the routing device is rejecting). • **Input SA rejects**—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • **Output packet count**—Number of packets that the filter has given to the MAC hardware. • **Output packet pad count**—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • **Output packet error count**—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment. • **CAM destination filters, CAM source filters**—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields should be 0.	extensive
PMA PHY	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information: • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. Any state other than **OK** indicates a problem.	extensive

 Copyright © 2018, Juniper Networks, Inc.

WIS section	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information: • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. Any state other than **OK** indicates a problem. Subfields are: • **BIP-B1**—Bit interleaved parity for SONET section overhead • **SEF**—Severely errored framing • **LOL**—Loss of light • **LOF**—Loss of frame • **ES-S**—Errored seconds (section) • S**ES-S**—Severely errored seconds (section) • **SEFS-S**—Severely errored framing seconds (section)	**extensive**
WIS line	(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **BIP-B2**—Bit interleaved parity for SONET line overhead • **REI-L**—Remote error indication (near-end line) • **RDI-L**—Remote defect indication (near-end line) • **AIS-L**—Alarm indication signal (near-end line) • **BERR-SF**—Bit error rate fault (signal failure) • **BERR-SD**—Bit error rate defect (signal degradation) • **ES-L**—Errored seconds (near-end line) • **SES-L**—Severely errored seconds (near-end line) • **UAS-L**—Unavailable seconds (near-end line) • **ES-LFE**—Errored seconds (far-end line) • **SES-LFE**—Severely errored seconds (far-end line) • **UAS-LFE**—Unavailable seconds (far-end line)	**extensive**

Copyright © 2018, Juniper Networks, Inc.

WIS path	(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. Any state other than **OK** indicates a problem. Subfields are: • **BIP-B3**—Bit interleaved parity for SONET section overhead • **REI-P**—Remote error indication • **LOP-P**—Loss of pointer (path) • **AIS-P**—Path alarm indication signal • **RDI-P**—Path remote defect indication • **UNEQ-P**—Path unequipped • **PLM-P**—Path payload label mismatch • **ES-P**—Errored seconds (near-end STS path) • **SES-P**—Severely errored seconds (near-end STS path) • **UAS-P**—Unavailable seconds (near-end STS path) • **SES-PFE**—Severely errored seconds (far-end STS path) • **UAS-PFE**—Unavailable seconds (far-end STS path)	**extensive**

Autonegotiation information	Information about link autonegotiation.
	• **Negotiation status:**
	  ◦ **Incomplete**—Ethernet interface has the speed or link mode configured.
	  ◦ **No autonegotiation**—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
	  ◦ **Complete**—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• **Link partner status**—**OK** when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• **Link partner:**
	  ◦ **Link mode**—Depending on the capability of the attached Ethernet device, either **Full-duplex** or **Half-duplex**.
	  ◦ **Flow control**—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is **None**. For Gigabit Ethernet interfaces, types are **Symmetric** (link partner supports **PAUSE** on receive and transmit), **Asymmetric** (link partner supports **PAUSE** on transmit), and **Symmetric/Asymmetric** (link partner supports both **PAUSE** on receive and transmit or only **PAUSE** receive).
	  ◦ **Remote fault**—Remote fault information from the link partner—**Failure** indicates a receive link error. **OK** indicates that the link partner is receiving. **Negotiation error** indicates a negotiation error. **Offline** indicates that the link partner is going offline.
	• **Local resolution**—Information from the link partner:
	  ◦ **Flow control**—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are **Symmetric** (link partner supports **PAUSE** on receive and transmit), **Asymmetric** (link partner supports **PAUSE** on transmit), and **Symmetric/Asymmetric** (link partner supports both **PAUSE** on receive and transmit or only **PAUSE** receive).
	  ◦ **Remote fault**—Remote fault information. **Link OK** (no error detected on receive), **Offline** (local interface is offline), and **Link Failure** (link error detected on receive).
		extensive

 Copyright © 2018, Juniper Networks, Inc.

Received path trace, Transmitted path trace	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the routing device at the other end of the fiber. The transmitted path trace value is the message that this routing device transmits.	**extensive**
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number.	**extensive**
CoS information	Information about the CoS queue for the physical interface. • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP interface index number for the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under “Common Output Fields Description” on page 1110.	All levels

VLAN-Tag	Rewrite profile applied to incoming or outgoing frames on the outer (**Out**) VLAN tag or for both the outer and inner (**In**) VLAN tags. • **push**—An outer VLAN tag is pushed in front of the existing VLAN tag. • **pop**—The outer VLAN tag of the incoming frame is removed. • **swap**—The outer VLAN tag of the incoming frame is overwritten with the user specified VLAN tag information. • **push**—An outer VLAN tag is pushed in front of the existing VLAN tag. • **push-push**—Two VLAN tags are pushed in from the incoming frame. • **swap-push**—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • **swap-swap**—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user specified VLAN tag value. • **pop-swap**—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • **pop-pop**—Both the outer and inner VLAN tags of the incoming frame are removed.	**brief detail extensive** none
Demux:	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: • Source Family Inet • Destination Family Inet	**detail extensive** none
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under “Common Output Fields Description” on page 1110.	**detail extensive** none
MTU	Maximum transmission unit size on the logical interface.	**detail extensive** none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	**detail extensive** none
Traffic statistics	Number and rate of bytes and packets received and transmitted on the specified interface set. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Input packets, Output packets**—Number of packets received and transmitted on the interface set.	**detail extensive**
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	**extensive**
Local statistics	Number and rate of bytes and packets destined to the routing device.	**extensive**

 Copyright © 2018, Juniper Networks, Inc.

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 1110.	**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 1110.	**detail extensive** none
protocol-family	Protocol family configured on the logical interface. If the protocol is **inet**, the IP address of the interface is also displayed.	**brief**
Flags	Information about address flag (possible values are described in the "Addresses Flags" section under "Common Output Fields Description" on page 1110.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interlace.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Copyright © 2018, Juniper Networks, Inc.

For Gigabit Ethernet IQ PICs, traffic and MAC statistics output varies. The following table 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). 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 68: Gigabit and 10 Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type

Interface Type	Sample Command	Byte and Octet Counts Include	Comments
Inbound physical interface	**show interfaces ge-0/3/0 extensive**	Traffic statistics: Input bytes: 496 bytes per packet, representing the Layer 2 packet MAC statistics: Received octets: 500 bytes per packet, representing the Layer 2 packet + 4 bytes	The additional 4 bytes are for the CRC.
Inbound logical interface	**show interfaces ge-0/3/0.50 extensive**	Traffic statistics: Input bytes: 478 bytes per packet, representing the Layer 3 packet	
Outbound physical interface	**show interfaces ge-0/0/0 extensive**	Traffic statistics: Input bytes: 490 bytes per packet, representing the Layer 3 packet + 12 bytes MAC statistics: Received octets: 478 bytes per packet, representing the Layer 3 packet	For input bytes, the additional 12 bytes 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	

Table 58 on page 1153 lists the output fields for the **show interfaces** command. Output fields are listed in the approximate order in which they appear.

 Copyright © 2018, Juniper Networks, Inc.

Table 69: show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface.	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
Link-level type	Encapsulation being used on the physical interface.	All levels
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
MTU	Maximum transmission unit size on the physical interface.	All levels
Link mode	Link mode: Full-duplex or Half-duplex.	
Speed	Speed at which the interface is running.	All levels
BPDU error	Bridge protocol data unit (BPDU) error: Detected or None	
Loopback	Loopback status: **Enabled** or **Disabled**. If loopback is enabled, type of loopback: **Local** or **Remote**.	All levels
Source filtering	Source filtering status: **Enabled** or **Disabled**.	All levels
Flow control	Flow control status: **Enabled** or **Disabled**.	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: **Enabled** or **Disabled**.	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: • **Online**—Autonegotiation is manually configured as online. • **Offline**—Autonegotiation is manually configured as offline.	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Link flags	Information about the physical link.	All levels
CoS queues	Number of CoS queues configured.	**detail extensive** none
Current address	Configured MAC address.	**detail extensive** none

Copyright © 2018, Juniper Networks, Inc.

Table 69: show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
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
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. These fields can contain the value **None** or **Link**. • **None**—There are no active defects or alarms. • **Link**—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	**detail extensive** none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 69: show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface. • **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** . • **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**
Output errors	Output errors on the interface. • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), 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; therefore, for Gigabit Ethernet PICs, this number must 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 must 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 interfaces. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 69: show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Ingress queues	Total number of ingress queues supported on the specified interface.	**extensive**
Queue counters and queue number	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**

MAC statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC subsystem, including the following:
	• **Total octets** and **total packets**—Total number of octets and packets.
	• **Unicast packets, Broadcast packets,** and **Multicast packets**—Number of unicast, broadcast, and multicast packets.
	• **CRC/Align errors**—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).
	• **FIFO error**—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning.
	• **MAC control frames**—Number of MAC control frames.
	• **MAC pause frames**—Number of MAC control frames with **pause** operational code.
	• **Oversized frames**—There are two possible conditions regarding the number of oversized frames:
	  ◦ Packet length exceeds 1518 octets, or
	  ◦ Packet length exceeds MRU
	• **Jabber frames**—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.
	• **Fragment frames**—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.
	• **VLAN tagged frames**—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not.
	• **Code violations**—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error."
		extensive

 Copyright © 2018, Juniper Networks, Inc.

Table 69: show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Filter statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected. • **Input packet count**—Number of packets received from the MAC hardware that the filter processed. • **Input packet rejects**—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • **Input DA rejects**—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the device from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local device (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 must be 0.	**extensive**

Autonegotiation information	Information about link autonegotiation.
	• **Negotiation status:**
	  ◦ **Incomplete**—Ethernet interface has the speed or link mode configured.
	  ◦ **No autonegotiation**—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
	  ◦ **Complete**—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
		extensive

Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number.	**extensive**

Table 69: show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface. • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**
Interface transmit statistics	Status of the **interface-transmit-statistics** configuration: Enabled or Disabled.	**detail extensive**
Queue counters (Egress)	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**
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.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Traffic statistics	Number and rate of bytes and packets received and transmitted on the specified interface set. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Input packets, Output packets**—Number of packets received and transmitted on the interface set.	**detail extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 69: show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Local statistics	Number and rate of bytes and packets destined to the device.	**extensive**
Transit statistics	Number and rate of bytes and packets transiting the switch. NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the **Output bytes** and **Output packets** interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.	**extensive**
Security	Security zones that interface belongs to.	**extensive**
Flow Input statistics	Statistics on packets received by flow module.	**extensive**
Flow Output statistics	Statistics on packets sent by flow module.	**extensive**
Flow error statistics (Packets dropped due to)	Statistics on errors in the flow module.	**extensive**
Protocol	Protocol family.	**detail extensive** none
MTU	Maximum transmission unit size on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Route table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive** none
Flags	Information about protocol family flags. .	**detail extensive**
Addresses, Flags	Information about the address flags..	**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**

Copyright © 2018, Juniper Networks, Inc.

Sample Output Gigabit Ethernet

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: 00:00:5e:00:53:7c, Hardware address: 00:00:5e:00:53:7c
  Last flapped   : 2006-08-10 17:25:10 PDT (00:01:08 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Ingress rate at Packet Forwarding Engine      : 0 bps (0 pps)
  Ingress drop rate at Packet Forwarding Engine : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None

  Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69)
    Flags: SNMP-Traps 0x4000
    VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push
    0x8100.512 0x8100.513)
    Encapsulation: VLAN-CCC
    Egress account overhead: 100
    Ingress account overhead: 90
    Input packets : 0
    Output packets: 0
    Protocol ccc, MTU: 1522
      Flags: Is-Primary
```

show interfaces (Gigabit Ethernet on MX Series Routers)

```
user@host> show interfaces ge-2/2/2
Physical interface: ge-2/2/2, Enabled, Physical link is Up
  Interface index: 156, SNMP ifIndex: 188
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, MAC-REWRITE Error: None,
 Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Schedulers     : 0
  Current address: 00:00:5e:00:53:c0, Hardware address: 00:00:5e:00:53:76
  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: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
    Input packets : 10232
    Output packets: 10294
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
```

 Copyright © 2018, Juniper Networks, Inc.

```
        Addresses, Flags: Is-Preferred Is-Primary
          Destination: 203.0.113/24, Local: 203.0.113.1, Broadcast: 203.0.113.255
      Protocol inet6, MTU: 1500
       Max nh cache: 4, New hold nh limit: 100000, Curr nh cnt: 4, Curr new hold
cnt: 4, NH drop cnt: 0
        Flags: Is-Primary
        Addresses, Flags: Is-Default Is-Preferred Is-Primary
          Destination: 2001:db8:/32, Local: 2001:db8::5
        Addresses, Flags: Is-Preferred
          Destination: 2001:db8:1::/32, Local: 2001:db8:223:9cff:fe9f:3e78
      Protocol multiservice, MTU: Unlimited
        Flags: Is-Primary
```

show interfaces (link degrade status)

```
user@host> show interfaces et-3/0/0
Physical interface: et-3/0/0, Enabled, Physical link is Down
  Interface index: 157, SNMP ifIndex: 537
  Link-level type: Ethernet, MTU: 1514, MRU: 0, Speed: 100Gbps, BPDU 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
  Current address: 54:e0:32:23:9d:38, Hardware address: 54:e0:32:23:9d:38
  Last flapped   : 2014-06-18 02:36:38 PDT (02:50:50 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0
  Link Degrade* :
  Link Monitoring                   :  Enable
  Link Degrade Set Threshold:       :  1E-7
  Link Degrade Clear Threshold:     :  1E-12
  Estimated BER                     :  1E-7
  Link-degrade event                :  Seconds     Count    State
                                           782         1     Defect Active
```

show interfaces extensive (Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration)

```
user@host> show interfaces ge-2/1/2 extensive | match "output|interface"
Physical interface: ge-2/1/2, Enabled, Physical link is Up
  Interface index: 151, SNMP ifIndex: 530, Generation: 154
  Interface flags: SNMP-Traps Internal: 0x4000
   Output bytes  :         240614363944           772721536 bps
   Output packets:           3538446506             1420444 pps
    Direction : Output
  Interface transmit statistics: Enabled

  Logical interface ge-2/1/2.0 (Index 331) (SNMP ifIndex 955) (Generation 146)
     Output bytes  :         195560312716           522726272 bps
     Output packets:           4251311146             1420451 pps

user@host> show interfaces ge-5/2/0.0 statistics detail
Logical interface ge-5/2/0.0 (Index 71) (SNMP ifIndex 573) (Generation 135)
    Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
```

Copyright © 2018, Juniper Networks, Inc.

```
    Egress account overhead: 100
    Ingress account overhead: 90
    Traffic statistics:
     Input  bytes  :              271524
     Output bytes  :            37769598
     Input  packets:                3664
     Output packets:              885790
     IPv6 transit statistics:
      Input  bytes  :                  0
      Output bytes  :           16681118
      Input  packets:                  0
      Output packets:             362633
    Local statistics:
     Input  bytes  :              271524
     Output bytes  :              308560
     Input  packets:                3664
     Output packets:                3659
    Transit statistics:
     Input  bytes  :                   0                    0 bps
     Output bytes  :            37461038                    0 bps
     Input  packets:                   0                    0 pps
     Output packets:              882131                    0 pps
     IPv6 transit statistics:
      Input  bytes  :                  0     0 bps
      Output bytes  :           16681118     0 bps
      Input  packets:                  0     0 pps
      Output packets:             362633     0 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
```

 Copyright © 2018, Juniper Networks, Inc.

```
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:7c, Hardware address: 00:00:5e:00:53:7c
  Last flapped   : 2006-08-09 17:17:00 PDT (01:31:33 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                  0                   0 bps
   Output bytes  :                  0                   0 bps
   Input  packets:                  0                   0 pps
   Output packets:                  0                   0 pps
  Ingress traffic statistics at Packet Forwarding Engine:
   Input  bytes  :                  0                   0 bps
   Input  packets:                  0                   0 pps
   Drop   bytes  :                  0                   0 bps
   Drop   packets:                  0                   0 pps
  Ingress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                  0                    0                    0

    1 expedited-fo                 0                    0                    0

    2 assured-forw                 0                    0                    0

    3 network-cont                 0                    0                    0

  Egress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                  0                    0                    0

    1 expedited-fo                 0                    0                    0

    2 assured-forw                 0                    0                    0

    3 network-cont                 0                    0                    0

  Active alarms  : None
  Active defects : None

  Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69) (Generation 140)
    Flags: SNMP-Traps 0x4000
    VLAN-Tag [0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530)
Out(swap-push 0x8100.512 0x8100.513)
    Encapsulation: VLAN-CCC
    Egress account overhead: 100
    Ingress account overhead: 90
    Traffic statistics:
     Input  bytes  :                  0
     Output bytes  :                  0
     Input  packets:                  0
     Output packets:                  0
    Local statistics:
     Input  bytes  :                  0
     Output bytes  :                  0
     Input  packets:                  0
     Output packets:                  0
    Transit statistics:
     Input  bytes  :                  0                   0 bps
     Output bytes  :                  0                   0 bps
     Input  packets:                  0                   0 pps
```

```
    Output packets:                   0                    0 pps
   Protocol ccc, MTU: 1522, Generation: 149, Route table: 0
     Flags: Is-Primary

 Logical interface ge-3/0/2.32767 (Index 71) (SNMP ifIndex 70)
 (Generation 139)
   Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
   Traffic statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
   Local statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
   Transit statistics:
    Input  bytes  :                   0                    0 bps
    Output bytes  :                   0                    0 bps
    Input  packets:                   0                    0 pps
    Output packets:                   0                    0 pps
```

show interfaces extensive (Gigabit Ethernet IQ2)

```
user@host> show interfaces ge-7/1/3 extensive
Physical interface: ge-7/1/3, Enabled, Physical link is Up
  Interface index: 170, SNMP ifIndex: 70, Generation: 171
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4004000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Schedulers     : 256
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:74, Hardware address: 00:00:5e:00:53:74
  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,
```

 Copyright © 2018, Juniper Networks, Inc.

```
  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Ingress queues: 4 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets

  0 best-effort           418390823            418390823                    0

  1 expedited-fo                  0                    0                    0

  2 assured-forw                  0                    0                    0

  3 network-cont               7133                 7133                    0

Egress queues: 4 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets

  0 best-effort                1031                 1031                    0

  1 expedited-fo                  0                    0                    0

  2 assured-forw                  0                    0                    0

  3 network-cont              77872                77872                    0

Active alarms  : None
Active defects : None
MAC statistics:                      Receive         Transmit
  Total octets                   38910844056          7174605
  Total packets                    418398473            78903
  Unicast packets               408021893366             1026
  Broadcast packets                       10               12
  Multicast packets                418398217            77865
  CRC/Align errors                         0                0
  FIFO errors                              0                0
  MAC control frames                       0                0
  MAC pause frames                         0                0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0  OTN Received Overhead Bytes:
  APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58
  Payload Type: 0x08
OTN Transmitted Overhead Bytes:
  APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
  Payload Type: 0x08
Filter statistics:
  Input packet count               418398473
  Input packet rejects                   479
  Input DA rejects                       479
  Input SA rejects                         0
  Output packet count                                   78903
  Output packet pad count                                   0
  Output packet error count                                 0
  CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Complete
  Link partner:
      Link mode: Full-duplex, Flow control: Symmetric/Asymmetric,
      Remote fault: OK
  Local resolution:
```

```
        Flow control: Symmetric, Remote fault: Link OK
  Packet Forwarding Engine configuration:
    Destination slot: 7
  CoS information:
    Direction : Output
    CoS transmit queue          Bandwidth          Buffer     Priority   Limit
                              %            bps     %           usec
    0 best-effort            95      950000000    95              0
low    none
    3 network-control         5       50000000     5              0
low    none
    Direction : Input
    CoS transmit queue          Bandwidth          Buffer     Priority   Limit
                              %            bps     %           usec
    0 best-effort            95      950000000    95              0
low    none
    3 network-control         5       50000000     5              0
low    none

  Logical interface ge-7/1/3.0 (Index 70) (SNMP ifIndex 85) (Generation 150)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :             812400
     Output bytes  :            1349206
     Input  packets:               9429
     Output packets:               9449
     IPv6 transit statistics:
      Input  bytes  :                 0
      Output bytes  :                 0
      Input  packets:                 0
      Output packets:                 0
    Local statistics:
     Input  bytes  :             812400
     Output bytes  :            1349206
     Input  packets:               9429
     Output packets:               9449
    Transit statistics:
     Input  bytes  :                  0                7440 bps
     Output bytes  :                  0                7888 bps
     Input  packets:                  0                  10 pps
     Output packets:                  0                  11 pps
     IPv6 transit statistics:
      Input  bytes  :                 0
      Output bytes  :                 0
      Input  packets:                 0
      Output packets:                 0
    Protocol inet, MTU: 1500, Generation: 169, Route table: 0
      Flags: Is-Primary, Mac-Validate-Strict
      Mac-Validate Failures: Packets: 0, Bytes: 0
      Addresses, Flags: Is-Preferred Is-Primary
      Input Filters: F1-ge-3/0/1.0-in, F3-ge-3/0/1.0-in
      Output Filters: F2-ge-3/0/1.0-out (53)
      Destination: 203.0.113/24, Local: 203.0.113.2, Broadcast: 203.0.113.255,
        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 56 on page 1123.

show interfaces (Gigabit Ethernet Unnumbered Interface)

```
user@host> show interfaces ge-3/2/0
Physical interface: ge-3/2/0, Enabled, Physical link is Up
  Interface index: 148, SNMP ifIndex: 50
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Current address: 00:00:5e:00:53:f8, Hardware address: 00:00:5e:00:53: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: 203.0.113.22
```

show interfaces (ACI Interface Set Configured)

```
user@host> show interfaces ge-1/0/0.4001
  Logical interface ge-1/0/0.4001 (Index 340) (SNMP ifIndex 548)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.4001 ]  Encapsulation: PPP-over-

    Ethernet
    ACI VLAN:
      Dynamic Profile: aci-vlan-set-profile
    PPPoE:
      Dynamic Profile: aci-vlan-pppoe-profile,
      Service Name Table: None,
      Max Sessions: 32000, Max Sessions VSA Ignore: Off,
      Duplicate Protection: On, Short Cycle Protection: Off,
      Direct Connect: Off,
      AC Name: nbc
    Input packets : 9
    Output packets: 8
    Protocol multiservice, MTU: Unlimited
```

show interfaces (ALI Interface Set)

```
user@host> show interfaces ge-1/0/0.10
  Logical interface ge-1/0/0.10 (Index 346) (SNMP ifIndex 554) (Generation 155)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.10 ]  Encapsulation: ENET2
    Line Identity:
```

Copyright © 2018, Juniper Networks, Inc.

```
      Dynamic Profile: ali-set-profile
      Circuit-id Remote-id Accept-no-ids
    PPPoE:
      Dynamic Profile: ali-vlan-pppoe-profile,
      Service Name Table: None,
      Max Sessions: 32000, Max Sessions VSA Ignore: Off,
      Duplicate Protection: On, Short Cycle Protection: Off,
      Direct Connect: Off,
      AC Name: nbc
    Input packets : 9
    Output packets: 8
    Protocol multiservice, MTU: Unlimited
```

Sample Output Gigabit Ethernet

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, IQ2)

```
user@host> show interfaces xe-5/0/0 extensive
Physical interface: xe-5/0/0, Enabled, Physical link is Up
  Interface index: 177, SNMP ifIndex: 99, Generation: 178
  Link-level type: Ethernet, MTU: 1518, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Enabled,
  Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Schedulers     : 1024
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:f6, Hardware address: 00:00:5e:00:53: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
```

 Copyright © 2018, Juniper Networks, Inc.

```
  2 assured-forw                 0                    0                    0

  3 network-cont                 0                    0                    0

Egress queues: 4 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets

  0 best-effort                  0                    0                    0

  1 expedited-fo                 0                    0                    0

  2 assured-forw                 0                    0                    0

  3 network-cont                 0                    0                    0

Active alarms  : None
Active defects : None
PCS statistics                     Seconds
  Bit errors                           0
  Errored blocks                       0
MAC statistics:                      Receive         Transmit
  Total octets                    6970332384                0
  Total packets                     81050506                0
  Unicast packets                   81050000                0
  Broadcast packets                      506                0
  Multicast packets                        0                0
  CRC/Align errors                         0                0
  FIFO errors                              0                0
  MAC control frames                       0                0
  MAC pause frames                         0                0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0
Filter statistics:
  Input packet count                81050506
  Input packet rejects                   506
  Input DA rejects                         0
  Input SA rejects                         0
  Output packet count                                       0
  Output packet pad count                                   0
  Output packet error count                                 0
  CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 5
CoS information:
  Direction : Output
  CoS transmit queue          Bandwidth              Buffer Priority   Limit
                       %            bps     %        usec
  0 best-effort       95      950000000    95         0          low     none
  3 network-control    5       50000000     5        0          low     none

  Direction : Input
  CoS transmit queue          Bandwidth              Buffer Priority   Limit
                       %            bps     %        usec
  0 best-effort       95      950000000    95         0          low     none
  3 network-control    5       50000000     5         0          low     none

Logical interface xe-5/0/0.0 (Index 71) (SNMP ifIndex 95) (Generation 195)
```

```
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ]  Encapsulation: ENET2
    Egress account overhead: 100
    Ingress account overhead: 90
    Traffic statistics:
     Input  bytes  :                  0
     Output bytes  :                 46
     Input  packets:                  0
     Output packets:                  1
     IPv6 transit statistics:
      Input  bytes  :                 0
      Output bytes  :                 0
      Input  packets:                 0
      Output packets:                 0
    Local statistics:
     Input  bytes  :                  0
     Output bytes  :                 46
     Input  packets:                  0
     Output packets:                  1
    Transit statistics:
     Input  bytes  :                  0                    0 bps
     Output bytes  :                  0                    0 bps
     Input  packets:                  0                    0 pps
     Output packets:                  0                    0 pps
     IPv6 transit statistics:
      Input  bytes  :                 0
      Output bytes  :                 0
      Input  packets:                 0
      Output packets:                 0
    Protocol inet, MTU: 1500, Generation: 253, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.0.2/24, Local: 192.0.2.1, Broadcast: 192.0.2.255,
Generation: 265
    Protocol multiservice, MTU: Unlimited, Generation: 254, Route table: 0
      Flags: None
      Policer: Input: __default_arp_policer__
```

show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode)

```
user@host> show interfaces xe-1/0/0 extensive
Physical interface: xe-1/0/0, Enabled, Physical link is Up
  Interface index: 141, SNMP ifIndex: 34, Generation: 47
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Loopback: Disabled
  WAN-PHY mode
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps 16384
  Link flags     : None
  CoS queues     : 4 supported
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:9d, Hardware address: 00:00:5e:00:53:9d
  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,
```

 Copyright © 2018, Juniper Networks, Inc.

```
    Resource errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0,
    Aged packets: 0, FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0,
    Resource errors: 0
  Queue counters:       Queued packets  Transmitted packets     Dropped packets
    0 best-effort                    0                    0                   0
    1 expedited-fo                   0                    0                   0
    2 assured-forw                   0                    0                   0
    3 network-cont                   0                    0                   0
  Active alarms : LOL, LOS, LBL
  Active defects: LOL, LOS, LBL, SEF, AIS-L, AIS-P
  PCS statistics             Seconds       Count
    Bit errors                     0           0
    Errored blocks                 0           0
  MAC statistics:                      Receive         Transmit
    Total octets                             0                0
    Total packets                            0                0
    Unicast packets                          0                0
    Broadcast packets                        0                0
    Multicast packets                        0                0
    CRC/Align errors                         0                0
    FIFO errors                              0                0
    MAC control frames                       0                0
    MAC pause frames                         0                0
    Oversized frames                         0
    Jabber frames                            0
    Fragment frames                          0
    VLAN tagged frames                       0
    Code violations                          0
  Filter statistics:
    Input packet count                       0
    Input packet rejects                     0
    Input DA rejects                         0
    Input SA rejects                         0
    Output packet count                                       0
    Output packet pad count                                   0
    Output packet error count                                 0
    CAM destination filters: 0, CAM source filters: 0
  PMA PHY:              Seconds        Count  State
    PLL lock                  0            0  OK
    PHY light             63159            1  Light Missing
 WIS section:
   BIP-B1                     0            0
   SEF                   434430       434438  Defect Active
   LOS                   434430            1  Defect Active
   LOF                   434430            1  Defect Active
   ES-S                  434430
   SES-S                 434430
   SEFS-S                434430
 WIS line:
   BIP-B2                     0            0
   REI-L                      0            0
   RDI-L                      0            0  OK
   AIS-L                 434430            1  Defect Active
   BERR-SF                    0            0  OK
   BERR-SD                    0            0  OK
   ES-L                  434430
   SES-L                 434430
   UAS-L                 434420
   ES-LFE                     0
```

```
  SES-LFE                  0
  UAS-LFE                  0
WIS path:
  BIP-B3                   0            0
  REI-P                    0            0
  LOP-P                    0            0  OK
  AIS-P               434430            1  Defect Active
  RDI-P                    0            0  OK
  UNEQ-P                   0            0  OK
  PLM-P                    0            0  OK
  ES-P                434430
  SES-P               434430
  UAS-P               434420
  ES-PFE                   0
  SES-PFE                  0
  UAS-PFE                  0
Received path trace:
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
Transmitted path trace: orissa so-1/0/0
  6f 72 69 73 73 61 20 73 6f 2d 31 2f 30 2f 30 00   orissa so-1/0/0.
  Packet Forwarding Engine configuration:
    Destination slot: 1
  CoS information:
    CoS transmit queue         Bandwidth          Buffer     Priority   Limit
                      %              bps      %   bytes
    0 best-effort      95      950000000     95       0          low     none
    3 network-control   5       50000000      5       0          low     none
```

show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC)

```
user@host> show interfaces ge-7/0/0 extensive
Physical interface: ge-7/0/0, Enabled, Physical link is Down
  Interface index: 143, SNMP ifIndex: 508, Generation: 208
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, BPDU Error: None,
  MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Enabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
  Link flags     : None
  Wavelength     : 1550.12 nm, Frequency: 193.40 THz
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:72, Hardware address: 00:00:5e:00:53:72
  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,
```

 Copyright © 2018, Juniper Networks, Inc.

```
  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets

  0 best-effort                    0                    0                    0

  1 expedited-fo                   0                    0                    0

  2 assured-forw                   0                    0                    0

  3 network-cont
Queue number:         Mapped forwarding classes
  0                   best-effort
  1                   expedited-forwarding
  2                   assured-forwarding
  3                   network-control
Active alarms  : LINK
Active defects : LINK
MAC statistics:                      Receive         Transmit
  Total octets                             0                0
  Total packets                            0                0
  Unicast packets                          0                0
  Broadcast packets                        0                0
  Multicast packets                        0                0
  CRC/Align errors                         0                0
  FIFO errors                              0                0
  MAC control frames                       0                0
  MAC pause frames                         0                0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0
  Total octets                             0                0
  Total packets                            0                0
  Unicast packets                          0                0
  Broadcast packets                        0                0
  Multicast packets                        0                0
  CRC/Align errors                         0                0
  FIFO errors                              0                0
  MAC control frames                       0                0
  MAC pause frames                         0                0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0
OTN alarms           :   None
OTN defects          :   None
OTN FEC Mode       : GFEC
OTN Rate               : Fixed Stuff Bytes 11.0957Gbps
OTN Line Loopback : Enabled
OTN FEC statistics :
  Corrected Errors                                      0
  Corrected Error Ratio (         0 sec average)     0e-0
OTN FEC alarms:          Seconds        Count  State
  FEC Degrade                  0            0  OK
  FEC Excessive                0            0  OK
OTN OC:                  Seconds        Count  State
  LOS                          2            1  OK
  LOF                      67164            2  Defect Active
```

```
    LOM                        67164           71  Defect Active
    Wavelength Lock                0            0  OK
  OTN OTU:
    AIS                            0            0  OK
    BDI                        65919         4814  Defect Active
    IAE                        67158            1  Defect Active
    TTIM                           7            1  OK
    SF                         67164            2  Defect Active
    SD                         67164            3  Defect Active
    TCA-ES                         0            0  OK
    TCA-SES                        0            0  OK
    TCA-UAS                       80           40  OK
    TCA-BBE                        0            0  OK
    BIP                            0            0  OK
    BBE                            0            0  OK
    ES                             0            0  OK
    SES                            0            0  OK
    UAS                          587            0  OK
    Received DAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Received SAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Transmitted DAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Transmitted SAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
  OTN Received Overhead Bytes:
    APS/PCC0: 0x02, APS/PCC1: 0x42, APS/PCC2: 0xa2, APS/PCC3: 0x48
    Payload Type: 0x03
  OTN Transmitted Overhead Bytes:
    APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
    Payload Type: 0x03
  Filter statistics:
    Input packet count                  0
    Input packet rejects                0
    Input DA rejects                    0
    Input SA rejects                    0
    Output packet count                                  0
    Output packet pad count                              0
    Output packet error count                            0
    CAM destination filters: 0, CAM source filters: 0
  Packet Forwarding Engine configuration:
    Destination slot: 7
  CoS information:
    Direction : Output
    CoS transmit queue              Bandwidth               Buffer Priority
Limit
                              %             bps     %            usec
    0 best-effort            95      9500000000    95               0       low
none
    3 network-control         5       500000000     5               0       low
none
    ...
```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode)

```
user@host> show interfaces xe-7/0/0 extensive
Physical interface: xe-7/0/0, Enabled, Physical link is Up
  Interface index: 173, SNMP ifIndex: 212, Generation: 174
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
Unidirectional: Enabled,
```

 Copyright © 2018, Juniper Networks, Inc.

```
  Loopback: None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
...
```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only)

```
user@host> show interfaces xe-7/0/0–tx extensive
Physical interface: xe-7/0/0-tx, Enabled, Physical link is Up
  Interface index: 176, SNMP ifIndex: 137, Generation: 177
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
Unidirectional: Tx-Only
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:83, Hardware address: 00:00:5e:00:53:83
  Last flapped   : 2007-06-01 09:08:19 PDT (3d 02:31 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :      322891152287160           9627472888 bps
   Input  packets:                    0                    0 pps
   Output packets:         328809727380              1225492 pps

...

  Filter statistics:
    Output packet count           328810554250
    Output packet pad count                  0
    Output packet error count                0
...

  Logical interface xe-7/0/0-tx.0 (Index 73) (SNMP ifIndex 138) (Generation 139)

    Flags: SNMP-Traps Encapsulation: ENET2
    Egress account overhead: 100
    Ingress account overhead: 90
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :      322891152287160
     Input  packets:                    0
     Output packets:         328809727380
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :      322891152287160           9627472888 bps
     Input  packets:                    0                    0 pps
     Output packets:         328809727380              1225492 pps
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
```

```
        Input  packets:                  0
        Output packets:                  0
      Protocol inet, MTU: 1500, Generation: 147, Route table: 0
        Addresses, Flags: Is-Preferred Is-Primary
          Destination: 10.11.12/24, Local: 10.11.12.13, Broadcast: 10.11.12.255,
Generation: 141
      Protocol multiservice, MTU: Unlimited, Generation: 148, Route table: 0
        Flags: None
        Policer: Input: __default_arp_policer__
```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only)

```
user@host> show interfaces xe-7/0/0–rx extensive
Physical interface: xe-7/0/0-rx, Enabled, Physical link is Up
  Interface index: 174, SNMP ifIndex: 118, Generation: 175
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
Unidirectional: Rx-Only
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:83, Hardware address: 00:00:5e:00:53:83
  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
```


Copyright © 2018, Juniper Networks, Inc.

```
     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: 192.0.2/24, Local: 192.0.2.1, Broadcast: 192.0.2.255,
Generation: 139
    Protocol multiservice, MTU: Unlimited, Generation: 146, Route table: 0
      Flags: None
      Policer: Input: __default_arp_policer__
```

Sample Output

Sample Output SRX Gigabit Ethernet

```
user@host> show interfaces ge-0/0/1
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, 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 Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:42 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK
  Interface transmit statistics: Disabled

  Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514)
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Security: Zone: public
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255
```

Sample Output SRX Gigabit Ethernet

```
user@host> show interfaces ge-0/0/1
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, 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 Down
```

```
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:42 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK
  Interface transmit statistics: Disabled

  Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514)
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Security: Zone: public
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255
```

show interfaces detail (Gigabit Ethernet)

user@host> show interfaces ge-0/0/1 detail

```
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510, Generation: 138
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, 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 Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w2d 00:00 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
  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  : LINK
  Active defects : LINK
```


Copyright © 2018, Juniper Networks, Inc.

```
Interface transmit statistics: Disabled

Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)
  Flags: Device-Down SNMP-Traps 0x0 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
  Security: Zone: public
  Flow Statistics :
  Flow Input statistics :
    Self packets :                     0
    ICMP packets :                     0
    VPN packets :                      0
    Multicast packets :                0
    Bytes permitted by policy :        0
    Connections established :          0
  Flow Output statistics:
    Multicast packets :                0
    Bytes permitted by policy :        0
  Flow error statistics (Packets dropped due to):
    Address spoofing:                  0
    Authentication failed:             0
    Incoming NAT errors:               0
    Invalid zone received packet:      0
    Multiple user authentications:     0
    Multiple incoming NAT:             0
    No parent for a gate:              0
    No one interested in self packets: 0
    No minor session:                  0
    No more sessions:                  0
    No NAT gate:                       0
    No route present:                  0
    No SA for incoming SPI:            0
    No tunnel found:                   0
    No session for a gate:             0
    No zone or NULL zone binding       0
    Policy denied:                     0
    Security association not active:   0
    TCP sequence number out of window: 0
    Syn-attack protection:             0
    User authentication errors:        0
  Protocol inet, MTU: 1500, Generation: 150, Route table: 0
    Flags: Sendbcast-pkt-to-re
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255, Generation:
150
```

Copyright © 2018, Juniper Networks, Inc.

show interfaces statistics st0.0 detail

```
user@host> show interfaces statistics st0.0 detail
  Logical interface st0.0 (Index 71) (SNMP ifIndex 609) (Generation 136)
    Flags: Up Point-To-Point SNMP-Traps Encapsulation: Secure-Tunnel
    Traffic statistics:
     Input  bytes  :        528152756774
     Output bytes  :        575950643520
     Input  packets:         11481581669
     Output packets:         12520666095
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :                   0            121859888 bps
     Output bytes  :                   0            128104112 bps
     Input  packets:                   0               331141 pps
     Output packets:                   0               348108 pps
    Security: Zone: untrust
    Allowed host-inbound traffic : any-service bfd bgp dvmrp igmp ldp msdp nhrp
ospf ospf3 pgm pim rip ripng router-discovery rsvp
    sap vrrp
    Flow Statistics :
    Flow Input statistics :
      Self packets :                    0
      ICMP packets :                    0
      VPN packets :                     0
      Multicast packets :               0
      Bytes permitted by policy :       525984295844
      Connections established :         7
    Flow Output statistics:
      Multicast packets :               0
      Bytes permitted by policy :       576003290222
    Flow error statistics (Packets dropped due to):
      Address spoofing:                 0
      Authentication failed:            0
      Incoming NAT errors:              0
      Invalid zone received packet:     0
      Multiple user authentications:    0
      Multiple incoming NAT:            0
      No parent for a gate:             0
      No one interested in self packets: 0
      No minor session:                 0
      No more sessions:                 0
      No NAT gate:                      0
      No route present:                 2000280
      No SA for incoming SPI:           0
      No tunnel found:                  0
      No session for a gate:            0
      No zone or NULL zone binding      0
      Policy denied:                    0
      Security association not active:  0
      TCP sequence number out of window: 0
      Syn-attack protection:            0
      User authentication errors:       0
    Protocol inet, MTU: 9192
    Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0,
 NH drop cnt: 0
    Generation: 155, Route table: 0
```


Copyright © 2018, Juniper Networks, Inc.

```
        Flags: Sendbcast-pkt-to-re
```

show interfaces extensive (Gigabit Ethernet)

```
user@host> show interfaces ge-0/0/1.0 extensive
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510, Generation: 138
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, 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 Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:57 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,
    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  : 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
```

Copyright © 2018, Juniper Networks, Inc.

```
    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: 2, CAM source filters: 0
  Autonegotiation information:
    Negotiation status: Incomplete
  Packet Forwarding Engine configuration:
    Destination slot: 0
  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
  Interface transmit statistics: Disabled

  Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)
    Flags: Device-Down SNMP-Traps 0x0 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
    Security: Zone: public
    Flow Statistics :
    Flow Input statistics :
      Self packets :                     0
      ICMP packets :                     0
      VPN packets :                      0
      Multicast packets :                0
      Bytes permitted by policy :        0
      Connections established :          0
    Flow Output statistics:
      Multicast packets :                0
      Bytes permitted by policy :        0
    Flow error statistics (Packets dropped due to):
      Address spoofing:                  0
```


Copyright © 2018, Juniper Networks, Inc.

```
  Authentication failed:            0
  Incoming NAT errors:              0
  Invalid zone received packet:     0
  Multiple user authentications:    0
  Multiple incoming NAT:            0
  No parent for a gate:             0
  No one interested in self packets: 0
  No minor session:                 0
  No more sessions:                 0
  No NAT gate:                      0
  No route present:                 0
  No SA for incoming SPI:           0
  No tunnel found:                  0
  No session for a gate:            0
  No zone or NULL zone binding      0
  Policy denied:                    0
  Security association not active:  0
  TCP sequence number out of window: 0
  Syn-attack protection:            0
  User authentication errors:       0
Protocol inet, MTU: 1500, Generation: 150, Route table: 0
  Flags: Sendbcast-pkt-to-re
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
    Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255,
    Generation: 150
```

show interfaces terse

```
user@host> show interfaces terse
    Interface               Admin Link Proto    Local                 Remote
    ge-0/0/0                up    up
    ge-0/0/0.0              up    up   inet     10.209.4.61/18
    gr-0/0/0                up    up
    ip-0/0/0                up    up
    st0                     up    up
    st0.1                   up    ready inet
    ls-0/0/0                up    up
    lt-0/0/0                up    up
    mt-0/0/0                up    up
    pd-0/0/0                up    up
    pe-0/0/0                up    up
    e3-1/0/0                up    up
    t3-2/0/0                up    up
    e1-3/0/0                up    up
    se-4/0/0                up    down
    t1-5/0/0                up    up
    br-6/0/0                up    up
    dc-6/0/0                up    up
    dc-6/0/0.32767          up    up
    bc-6/0/0:1              down  up
    bc-6/0/0:1.0            up    down
    dl0                     up    up
    dl0.0                   up    up   inet
    dsc                     up    up
    gre                     up    up
    ipip                    up    up
    lo0                     up    up
    lo0.16385               up    up   inet     10.0.0.1            --> 0/0
                                                10.0.0.16           --> 0/0
```

```
    lsi                     up     up
    mtun                    up     up
    pimd                    up     up
    pime                    up     up
    pp0                     up     up
```

show interfaces controller (Channelized E1 IQ with Logical E1)

```
user@host> show interfaces controller ce1-1/2/6

Controller                                                         Admin Link
ce1-1/2/6                                                          up    up
e1-1/2/6                                                           up    up
```

show interfaces controller (Channelized E1 IQ with Logical DS0)

```
user@host> show interfaces controller ce1-1/2/3

Controller                                                         Admin Link
ce1-1/2/3                                                          up    up

ds-1/2/3:1                                                         up    up

ds-1/2/3:2                                                         up    up
```

show interfaces descriptions

```
user@host> show interfaces descriptions
 Interface       Admin Link Description
 so-1/0/0        up    up   M20-3#1
 so-2/0/0        up    up   GSR-12#1
 ge-3/0/0        up    up   SMB-OSPF_Area300
 so-3/3/0        up    up   GSR-13#1
 so-3/3/1        up    up   GSR-13#2
 ge-4/0/0        up    up   T320-7#1
 ge-5/0/0        up    up   T320-7#2
 so-7/1/0        up    up   M160-6#1
 ge-8/0/0        up    up   T320-7#3
 ge-9/0/0        up    up   T320-7#4
 so-10/0/0       up    up   M160-6#2
 so-13/0/0       up    up   M20-3#2
 so-14/0/0       up    up   GSR-12#2
 ge-15/0/0       up    up   SMB-OSPF_Area100
 ge-15/0/1       up    up   GSR-13#3
```

show interfaces destination-class all

```
user@host> show interfaces destination-class all
Logical interface so-4/0/0.0
                                            Packets                 Bytes
      Destination class              (packet-per-second)    (bits-per-second)
                         gold                        0                     0
                                    (                0)  (                 0)
                       silver                        0                     0
                                    (                0)  (                 0)
  Logical interface so-0/1/3.0
                                            Packets                 Bytes
      Destination class              (packet-per-second)    (bits-per-second)
                         gold                        0                     0
                                    (                0)  (                 0)
```


Copyright © 2018, Juniper Networks, Inc.

```
                silver                          0                    0
                           (                    0) (                 0)
```

show interfaces diagnostics optics

```
user@host> show interfaces diagnostics optics ge-2/0/0
Physical interface: ge-2/0/0
    Laser bias current                        :  7.408 mA
    Laser output power                        :  0.3500 mW / -4.56 dBm
    Module temperature                        :  23 degrees C / 73 degrees F
    Module voltage                            :  3.3450 V
    Receiver signal average optical power     :  0.0002 mW / -36.99 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  On
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  On
    Laser bias current high alarm threshold   :  17.000 mA
    Laser bias current low alarm threshold    :  1.000 mA
    Laser bias current high warning threshold :  14.000 mA
    Laser bias current low warning threshold  :  2.000 mA
    Laser output power high alarm threshold   :  0.6310 mW / -2.00 dBm
    Laser output power low alarm threshold    :  0.0670 mW / -11.74 dBm
    Laser output power high warning threshold :  0.6310 mW / -2.00 dBm
    Laser output power low warning threshold  :  0.0790 mW / -11.02 dBm
    Module temperature high alarm threshold   :  95 degrees C / 203 degrees F
    Module temperature low alarm threshold    :  -25 degrees C / -13 degrees F
    Module temperature high warning threshold :  90 degrees C / 194 degrees F
    Module temperature low warning threshold  :  -20 degrees C / -4 degrees F
    Module voltage high alarm threshold       :  3.900 V
    Module voltage low alarm threshold        :  2.700 V
    Module voltage high warning threshold     :  3.700 V
    Module voltage low warning threshold      :  2.900 V
    Laser rx power high alarm threshold       :  1.2590 mW / 1.00 dBm
    Laser rx power low alarm threshold        :  0.0100 mW / -20.00 dBm
    Laser rx power high warning threshold     :  0.7940 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.0158 mW / -18.01 dBm
```

show interfaces far-end-interval coc12-5/2/0

```
user@host> show interfaces far-end-interval coc12-5/2/0
Physical interface: coc12-5/2/0, SNMP ifIndex: 121
  05:30-current:
    ES-L: 1, SES-L: 1, UAS-L: 0
```

```
  05:15-05:30:
    ES-L: 0, SES-L: 0, UAS-L: 0
  05:00-05:15:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:45-05:00:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:30-04:45:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:15-04:30:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:00-04:15:
...
```

show interfaces far-end-interval coc1-5/2/1:1

```
user@host> run show interfaces far-end-interval coc1-5/2/1:1
 Physical interface: coc1-5/2/1:1, SNMP ifIndex: 342
  05:30-current:
    ES-L: 1, SES-L: 1, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  05:15-05:30:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  05:00-05:15:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:45-05:00:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:30-04:45:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:15-04:30:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:00-04:15:
```

show interfaces filters

```
user@host> show interfaces filters
Interface       Admin Link Proto Input Filter         Output Filter
ge-0/0/0        up    up
ge-0/0/0.0      up    up   inet
                           iso
ge-5/0/0        up    up
ge-5/0/0.0      up    up   any                        f-any
                           inet                       f-inet
                           multiservice
gr-0/3/0        up    up
ip-0/3/0        up    up
mt-0/3/0        up    up
pd-0/3/0        up    up
pe-0/3/0        up    up
vt-0/3/0        up    up
at-1/0/0        up    up
at-1/0/0.0      up    up   inet
                           iso
at-1/1/0        up    down
at-1/1/0.0      up    down inet
                           iso
....
```

show interfaces flow-statistics (Gigabit Ethernet)

```
user@host> show interfaces flow-statistics ge-0/0/1.0
```

 Copyright © 2018, Juniper Networks, Inc.

```
  Logical interface ge-0/0/1.0 (Index 70) (SNMP ifIndex 49)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 5161
    Output packets: 83
    Security: Zone: zone2
    Allowed host-inbound traffic : bootp bfd bgp  dns dvmrp  ldp msdp nhrp ospf
pgm
    pim rip router-discovery rsvp sap vrrp dhcp finger ftp tftp ident-reset http
 https ike
   netconf ping rlogin rpm rsh snmp snmp-trap ssh telnet traceroute xnm-clear-text
 xnm-ssl
    lsping
    Flow Statistics :
    Flow Input statistics :
      Self packets :                     0
      ICMP packets :                     0
      VPN packets :                      2564
      Bytes permitted by policy :        3478
      Connections established :          1
    Flow Output statistics:
      Multicast packets :                0
      Bytes permitted by policy :        16994
    Flow error statistics (Packets dropped due to):
      Address spoofing:                  0
      Authentication failed:             0
      Incoming NAT errors:               0
      Invalid zone received packet:      0
      Multiple user authentications:     0
      Multiple incoming NAT:             0
      No parent for a gate:              0
      No one interested in self packets: 0
      No minor session:                  0
      No more sessions:                  0
      No NAT gate:                       0
      No route present:                  0
      No SA for incoming SPI:            0
      No tunnel found:                   0
      No session for a gate:             0
      No zone or NULL zone binding       0
      Policy denied:                     0
      Security association not active:   0
      TCP sequence number out of window: 0
      Syn-attack protection:             0
      User authentication errors:        0
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination:  203.0.113.1/24, Local: 203.0.113.2, Broadcast: 2.2.2.255
```

show interfaces interval (Channelized OC12)

```
user@host> show interfaces interval t3-0/3/0:0
Physical interface: t3-0/3/0:0, SNMP ifIndex: 23
  17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:28-17:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:13-17:28:
```

```
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:58-17:13:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
...
Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
    CES: 230, CSES: 230, SEFS: 230, UAS: 238
```

show interfaces interval (E3)

```
user@host> show interfaces interval e3-0/3/0
Physical interface: e3-0/3/0, SNMP ifIndex: 23
  17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:28-17:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:13-17:28:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:58-17:13:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
....
Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
    CES: 230, CSES: 230, SEFS: 230, UAS: 238
```

show interfaces interval (SONET/SDH) (SRX devices)

```
user@host> show interfaces interval so-0/1/0
Physical interface: so-0/1/0, SNMP ifIndex: 19
  20:02-current:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0,
    SES-P: 0, UAS-P: 0
  19:47-20:02:
    ES-S: 267, SES-S: 267, SEFS-S: 267, ES-L: 267, SES-L: 267, UAS-L: 267,
    ES-P: 267, SES-P: 267, UAS-P: 267
  19:32-19:47:
    ES-S: 56, SES-S: 56, SEFS-S: 56, ES-L: 56, SES-L: 56, UAS-L: 46, ES-P: 56,
    SES-P: 56, UAS-P: 46
  19:17-19:32:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0,
    SES-P: 0, UAS-P: 0
  19:02-19:17:
.....
```

show interfaces load-balancing (SRX devices)

```
user@host> show interfaces load-balancing
Interface  State           Last change  Member count
ams0       Up              1d 00:50     2
ams1       Up              00:00:59     2
```

 Copyright © 2018, Juniper Networks, Inc.

show interfaces load-balancing detail (SRX devices)

```
user@host>show interfaces load-balancing detail
Load-balancing interfaces detail
Interface        : ams0
  State          : Up
  Last change    : 1d 00:51
  Member count   : 2
  Members        :
      Interface     Weight   State
      mams-2/0/0     10        Active
      mams-2/1/0     10        Active
```

show interfaces mac-database (All MAC Addresses on a Port SRX devices)

```
user@host> show interfaces mac-database xe-0/3/3
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address        Input frames     Input bytes  Output frames   Output bytes
  00:00:00:00:00:00             1              56              0              0
  00:00:c0:01:01:02       7023810       323095260              0              0
  00:00:c0:01:01:03       7023810       323095260              0              0
  00:00:c0:01:01:04       7023810       323095260              0              0
  00:00:c0:01:01:05       7023810       323095260              0              0
  00:00:c0:01:01:06       7023810       323095260              0              0
  00:00:c0:01:01:07       7023810       323095260              0              0
  00:00:c0:01:01:08       7023809       323095214              0              0
  00:00:c0:01:01:09       7023809       323095214              0              0
  00:00:c0:01:01:0a       7023809       323095214              0              0
  00:00:c0:01:01:0b       7023809       323095214              0              0
  00:00:c8:01:01:02      30424784      1399540064       37448598     1722635508
  00:00:c8:01:01:03      30424784      1399540064       37448598     1722635508
  00:00:c8:01:01:04      30424716      1399536936       37448523     1722632058
  00:00:c8:01:01:05      30424789      1399540294       37448598     1722635508
  00:00:c8:01:01:06      30424788      1399540248       37448597     1722635462
  00:00:c8:01:01:07      30424783      1399540018       37448597     1722635462
  00:00:c8:01:01:08      30424783      1399540018       37448596     1722635416
  00:00:c8:01:01:09       8836796       406492616        8836795      406492570
  00:00:c8:01:01:0a      30424712      1399536752       37448521     1722631966
  00:00:c8:01:01:0b      30424715      1399536890       37448523     1722632058
  Number of MAC addresses : 21
```

show interfaces mac-database (All MAC Addresses on a Service SRX devices)

```
user@host> show interfaces mac-database xe-0/3/3
Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address        Input frames     Input bytes  Output frames   Output bytes
  00:00:00:00:00:00             1              56              0              0
```

```
00:00:c0:01:01:02      7023810    323095260          0           0
00:00:c0:01:01:03      7023810    323095260          0           0
00:00:c0:01:01:04      7023810    323095260          0           0
00:00:c0:01:01:05      7023810    323095260          0           0
00:00:c0:01:01:06      7023810    323095260          0           0
00:00:c0:01:01:07      7023810    323095260          0           0
00:00:c0:01:01:08      7023809    323095214          0           0
00:00:c0:01:01:09      7023809    323095214          0           0
00:00:c0:01:01:0a      7023809    323095214          0           0
00:00:c0:01:01:0b      7023809    323095214          0           0
00:00:c8:01:01:02     31016568   1426762128   38040381  1749857526
00:00:c8:01:01:03     31016568   1426762128   38040382  1749857572
00:00:c8:01:01:04     31016499   1426758954   38040306  1749854076
00:00:c8:01:01:05     31016573   1426762358   38040381  1749857526
00:00:c8:01:01:06     31016573   1426762358   38040381  1749857526
00:00:c8:01:01:07     31016567   1426762082   38040380  1749857480
00:00:c8:01:01:08     31016567   1426762082   38040379  1749857434
00:00:c8:01:01:09      9428580    433714680    9428580   433714680
00:00:c8:01:01:0a     31016496   1426758816   38040304  1749853984
00:00:c8:01:01:0b     31016498   1426758908   38040307  1749854122
```

show interfaces mac-database mac-address

```
user@host> show interfaces mac-database xe-0/3/3 mac-address (SRX devices)
00:00:c8:01:01:09
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address: 00:00:c8:01:01:09, Type: Configured,
    Input bytes      : 202324652
    Output bytes     : 202324560
    Input frames     : 4398362
    Output frames    : 4398360
  Policer statistics:
  Policer type       Discarded frames    Discarded bytes
  Output aggregate            3992386          183649756
```

show interfaces mc-ae (SRX devices)

```
user@host> show interfaces mc-ae ae0 unit 512
  Member Links    : ae0
  Local Status    : active
  Peer Status     : active
  Logical Interface        : ae0.512
  Core Facing Interface : Label Ethernet Interface
  ICL-PL          : Label Ethernet Interface
```

show interfaces media (SONET/SDH)

The following example displays the output fields unique to the **show interfaces media** command for a SONET interface (with no level of output specified):

 Copyright © 2018, Juniper Networks, Inc.

```
user@host> show interfaces media so-4/1/2
Physical interface: so-4/1/2, Enabled, Physical link is Up
  Interface index: 168, SNMP ifIndex: 495
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 16384
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 1783 (00:00:00 ago), Output: 1786 (00:00:08 ago)
  LCP state: Opened
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
  mpls: Not-configured
  CHAP state: Not-configured
  CoS queues     : 8 supported
  Last flapped   : 2005-06-15 12:14:59 PDT (04:31:29 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : None
  SONET defects  : None
  SONET errors:
     BIP-B1: 121, BIP-B2: 916, REI-L: 0, BIP-B3: 137, REI-P: 16747, BIP-BIP2: 0
  Received path trace: routerb so-1/1/2
  Transmitted path trace: routera so-4/1/2
```

show interfaces policers (SRX devices)

```
user@host> show interfaces policers
Interface       Admin Link Proto Input Policer         Output Policer
ge-0/0/0        up    up
ge-0/0/0.0      up    up   inet
                           iso
gr-0/3/0        up    up
ip-0/3/0        up    up
mt-0/3/0        up    up
pd-0/3/0        up    up
pe-0/3/0        up    up
...
so-2/0/0        up    up
so-2/0/0.0      up    up   inet  so-2/0/0.0-in-policer so-2/0/0.0-out-policer
                           iso
so-2/1/0        up    down
...
```

show interfaces policers interface-name (SRX devices)

```
user@host> show interfaces policers so-2/1/0
Interface       Admin Link Proto Input Policer         Output Policer
so-2/1/0        up    down
so-2/1/0.0      up    down inet  so-2/1/0.0-in-policer so-2/1/0.0-out-policer
                           iso
                           inet6
```

show interfaces queue (SRX devices)

The following truncated example shows the CoS queue sizes for queues 0, 1, and 3. Queue 1 has a queue buffer size (guaranteed allocated memory) of 9192 bytes.

Copyright © 2018, Juniper Networks, Inc.

```
user@host> show interfaces queue
Physical interface: ge-0/0/0, Enabled, Physical link is Up
  Interface index: 134, SNMP ifIndex: 509
Forwarding classes: 8 supported, 8 in use
Egress queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: class0
  Queued:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
  Transmitted:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
    Tail-dropped packets :                     0                     0 pps
    RL-dropped packets   :                     0                     0 pps
    RL-dropped bytes     :                     0                     0 bps
    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 Buffer Usage:
    Reserved buffer      :             118750000 bytes
  Queue-depth bytes      :
    Current              :                     0
..
..
Queue: 1, Forwarding classes: class1
..
..
 Queue Buffer Usage:
    Reserved buffer      :                  9192 bytes
  Queue-depth bytes      :
    Current              :                     0
..
..
Queue: 3, Forwarding classes: class3
  Queued:
..
..
Queue Buffer Usage:
    Reserved buffer      :               6250000 bytes
  Queue-depth bytes      :
    Current              :                     0
..
..
```

show interfaces redundancy (SRX devices)

```
user@host> show interfaces redundancy
Interface  State         Last change  Primary    Secondary  Current status
rsp0       Not present                sp-1/0/0   sp-0/2/0   both down
rsp1       On secondary  1d 23:56     sp-1/2/0   sp-0/3/0   primary down
rsp2       On primary    10:10:27     sp-1/3/0   sp-0/2/0   secondary down
rlsq0      On primary    00:06:24     lsq-0/3/0  lsq-1/0/0  both up
```

 Copyright © 2018, Juniper Networks, Inc.

show interfaces redundancy (Aggregated Ethernet SRX devices)

```
user@host> show interfaces redundancy
Interface  State          Last change  Primary     Secondary   Current status
rlsq0      On secondary    00:56:12     lsq-4/0/0   lsq-3/0/0   both up

ae0
ae1
ae2
ae3
ae4
```

show interfaces redundancy detail (SRX devices)

```
user@host> show interfaces redundancy detail
Interface        : rlsq0
  State          : On primary
  Last change    : 00:45:47
  Primary        : lsq-0/2/0
  Secondary      : lsq-1/2/0
  Current status : both up
  Mode           : hot-standby

Interface        : rlsq0:0
  State          : On primary
  Last change    : 00:45:46
  Primary        : lsq-0/2/0:0
  Secondary      : lsq-1/2/0:0
  Current status : both up
  Mode           : warm-standby
```

show interfaces routing brief (SRX devices)

```
user@host> show interfaces routing brief
Interface        State Addresses
so-5/0/3.0       Down  ISO   enabled
so-5/0/2.0       Up    MPLS  enabled
                       ISO   enabled
                       INET  192.168.2.120
                       INET  enabled
so-5/0/1.0       Up    MPLS  enabled
                       ISO   enabled
                       INET  192.168.2.130
                       INET  enabled
at-1/0/0.3       Up    CCC   enabled
at-1/0/0.2       Up    CCC   enabled
at-1/0/0.0       Up    ISO   enabled
                       INET  192.168.90.10
                       INET  enabled
lo0.0            Up    ISO   47.0005.80ff.f800.0000.0108.0001.1921.6800.5061.00
                       ISO   enabled
                       INET  127.0.0.1
fxp1.0           Up
fxp0.0           Up    INET  192.168.6.90
```

show interfaces routing detail (SRX devices)

```
user@host> show interfaces routing detail
so-5/0/3.0
  Index: 15, Refcount: 2, State: Up <Broadcast PointToPoint Multicast> Change:<>
```

```
  Metric: 0, Up/down transitions: 0, Full-duplex
  Link layer: HDLC serial line Encapsulation: PPP Bandwidth: 155Mbps
  ISO address (null)
    State: <Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
so-5/0/2.0
  Index: 14, Refcount: 7, State: <Up Broadcast PointToPoint Multicast> Change:<>

  Metric: 0, Up/down transitions: 0, Full-duplex
  Link layer: HDLC serial line Encapsulation: PPP Bandwidth: 155Mbps
  MPLS address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4458 bytes
  ISO address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
  INET address 192.168.2.120
    State: <Up Broadcast PointToPoint Multicast Localup> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
    Local address: 192.168.2.120
    Destination: 192.168.2.110/32
  INET address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
...
```

show interfaces routing-instance all (SRX devices)

```
user@host> show interfaces terse routing-instance all
Interface    Admin    Link   Proto    Local             Remote Instance
at-0/0/1     up       up     inet     10.0.0.1/24
ge-0/0/0.0   up       up     inet      192.168.4.28/24          sample-a
at-0/1/0.0   up       up     inet6     fe80::a:0:0:4/64         sample-b
so-0/0/0.0   up       up     inet     10.0.0.1/32
```

show interfaces snmp-index (SRX devices)

```
user@host> show interfaces snmp-index 33
Physical interface: so-2/1/1, Enabled, Physical link is Down
  Interface index: 149, SNMP ifIndex: 33
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down Point-To-Point SNMP-Traps 16384
  Link flags     : Keepalives
  CoS queues     : 8 supported
  Last flapped   : 2005-06-15 11:45:57 PDT (05:38:43 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : LOL, PLL, LOS
  SONET defects  : LOL, PLL, LOF, LOS, SEF, AIS-L, AIS-P
```

show interfaces source-class all (SRX devices)

```
user@host> show interfaces source-class all
Logical interface so-0/1/0.0
                                               Packets               Bytes
      Source class                 (packet-per-second)   (bits-per-second)
                             gold              1928095           161959980
```

 Copyright © 2018, Juniper Networks, Inc.

```
                                    (                889) (            597762)
                             bronze                    0                    0
                                    (                  0) (                 0)
                             silver                    0                    0
                                    (                  0) (                 0)
  Logical interface so-0/1/3.0
                                                 Packets                Bytes
      Source class                 (packet-per-second)    (bits-per-second)
                               gold                    0                    0
                                    (                  0) (                 0)
                             bronze                    0                    0
                                    (                  0) (                 0)
                             silver               116113              9753492
                                    (                939) (            631616)
```

show interfaces statistics (Fast Ethernet SRX devices)

```
user@host> show interfaces fe-1/3/1 statistics
Physical interface: fe-1/3/1, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 1042
  Description: ford fe-1/3/1
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:90:69:93:04:dc, Hardware address: 00:90:69:93:04:dc
  Last flapped   : 2006-04-18 03:08:59 PDT (00:01:24 ago)
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 0, Output errors: 0
  Active alarms  : None
  Active defects : None
  Logical interface fe-1/3/1.0 (Index 69) (SNMP ifIndex 50)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500
      Flags: Is-Primary, DCU, SCU-in
                                                 Packets                Bytes
      Destination class            (packet-per-second)    (bits-per-second)
                            silver1                    0                    0
                                    (                  0) (                 0)
                            silver2                    0                    0
                                    (                  0) (                 0)
                            silver3                    0                    0
                                    (                  0) (                 0)
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 10.27.245/24, Local: 10.27.245.2,
        Broadcast: 10.27.245.255
    Protocol iso, MTU: 1497
      Flags: Is-Primary
```

show interfaces switch-port (SRX devices)

```
user@host# show interfaces ge-slot/0/0 switch-port port-number
Port 0, Physical link is Up
        Speed: 100mbps, Auto-negotiation: Enabled
   Statistics:                          Receive         Transmit
          Total bytes                   28437086        21792250
          Total packets                   409145           88008
```

```
            Unicast packets                         9987              83817
            Multicast packets                     145002                  0
            Broadcast packets                     254156               4191
            Multiple collisions                       23                 10
            FIFO/CRC/Align errors                      0                  0
            MAC pause frames                           0                  0
            Oversized frames                           0
            Runt frames                                0
            Jabber frames                              0
            Fragment frames                            0
            Discarded frames                           0
          Autonegotiation information:
            Negotiation status: Complete
            Link partner:
                Link mode: Full-duplex, Flow control: None, Remote fault: OK, Link
   partner Speed: 100 Mbps
            Local resolution:
                Flow control: None, Remote fault: Link OK
```

show interfaces transport pm (SRX devices)

```
user@host> show interfaces transport pm all current et-0/1/0
Physical interface: et-0/1/0, SNMP ifIndex 515
  14:45-current           Elapse time:900 Seconds
  Near End       Suspect Flag:False             Reason:None
    PM                 COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED

  OTU-BBE               0                800             No              No
  OTU-ES                0                135             No              No
  OTU-SES               0                90              No              No
  OTU-UAS               427              90              No              No
  Far End        Suspect Flag:True              Reason:Unknown
    PM                 COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED

  OTU-BBE               0                800             No              No
  OTU-ES                0                135             No              No
  OTU-SES               0                90              No              No
  OTU-UAS               0                90              No              No
  Near End       Suspect Flag:False             Reason:None
    PM                 COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED

  ODU-BBE               0                800             No              No
  ODU-ES                0                135             No              No
  ODU-SES               0                90              No              No
  ODU-UAS               427              90              No              No
  Far End        Suspect Flag:True              Reason:Unknown
    PM                 COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED

  ODU-BBE               0                800             No              No
  ODU-ES                0                135             No              No
  ODU-SES               0                90              No              No
  ODU-UAS               0                90              No              No
  FEC            Suspect Flag:False             Reason:None
    PM                 COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED

  FEC-CorrectedErr        2008544300      0              NA              NA
  FEC-UncorrectedWords    0               0              NA              NA
  BER            Suspect Flag:False             Reason:None
    PM                  MIN       MAX       AVG        THRESHOLD      TCA-ENABLED
     TCA-RAISED
  BER                   3.6e-5    5.8e-5    3.6e-5     10.0e-3         No
```

 Copyright © 2018, Juniper Networks, Inc.

```
       Yes
Physical interface: et-0/1/0, SNMP ifIndex 515
  14:45-current
  Suspect Flag:True             Reason:Object Disabled
    PM                           CURRENT   MIN       MAX      AVG       THRESHOLD
      TCA-ENABLED       TCA-RAISED
                                                                           (MIN)
(MAX)     (MIN)  (MAX)        (MIN)  (MAX)
  Lane chromatic dispersion       0        0         0         0         0
 0         NA     NA          NA     NA
  Lane differential group delay   0        0         0         0         0
 0         NA     NA          NA     NA
    q Value                         120      120       120       120       0
   0         NA     NA          NA     NA
  SNR                              28       28        29        28        0
 0         NA     NA          NA     NA
  Tx output power(0.01dBm)        -5000    -5000     -5000     -5000     -300
 -100      No     No          No     No
  Rx input power(0.01dBm)         -3642    -3665     -3626     -3637     -1800
 -500      No     No          No     No
  Module temperature(Celsius)      46       46        46        46        -5
 75        No     No          No     No
  Tx laser bias current(0.1mA)    0        0         0         0         0
 0         NA     NA          NA     NA
  Rx laser bias current(0.1mA)    1270     1270      1270      1270      0
 0         NA     NA          NA     NA
  Carrier frequency offset(MHz)   -186     -186      -186      -186      -5000
 5000      No     No          No     No
```

show security zones (SRX devices)

```
user@host> show security zones
Functional zone: management
   Description: This is the management zone.
   Policy configurable: No
   Interfaces bound: 1
   Interfaces:
     ge-0/0/0.0
Security zone: Host
   Description: This is the host zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     fxp0.0
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
Security zone: def
   Description: This is the def zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/2.0
```

Copyright © 2018, Juniper Networks, Inc.

show interfaces (M Series, MX Series, T Series Routers, and PTX Series Management and Internal Ethernet)

List of Syntax

Syntax (M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface) on page 1422
Syntax (M Series, MX Series, T Series, and PTX Series Routers Internal Ethernet Interface) on page 1422

Syntax (M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface)

show interfaces em0 | fxp0
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Syntax (M Series, MX Series, T Series, and PTX Series Routers Internal Ethernet Interface)

show interfaces bcm0 | em0 | em1| fxp1 | fxp2 | ixgbe0 | ixgbe1
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced on PTX Series Packet Transport Routers for Junos OS Release 12.1.

Description

(M Series, T Series, TX Matrix Plus, and PTX Series devices only) Display status information about the management Ethernet and internal Ethernet interfaces.

Options

em0 | fxp0—(M Series, MX Series, T Series, and PTX Series) Display standard information about the management Ethernet interface. For supported Ethernet interface by chassis and Routing Engine, see “Supported Routing Engines by Router” on page 17.

bcm0 | em0 | em1 | fxp1 | fxp2 | ixgbe0 | ixgbe1—(M Series, MX Series, T Series, and PTX Series) Display standard information about the internal Ethernet interfaces. See “Supported Routing Engines by Router” on page 17 for the internal Ethernet interface names for each Routing Engine by hardware platform.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

 Copyright © 2018, Juniper Networks, Inc.

Required Privilege Level view

List of Sample Output

show interfaces brief (Management Ethernet) on page 1426
show interfaces (Management Ethernet) on page 1426
show interfaces (Management Ethernet [TX Matrix Plus Router]) on page 1427
show interfaces (Management Ethernet [PTX Series Packet Transport Routers]) on page 1427
show interfaces detail (Management Ethernet) on page 1427
show interfaces detail (Management Ethernet [TX Matrix Plus Router]) on page 1428
show interfaces detail (Management Ethernet [PTX Packet Transport Routers]) on page 1429
show interfaces extensive (Management Ethernet) on page 1429
show interfaces extensive (Management Ethernet [TX Matrix Plus Router]) on page 1430
show interfaces extensive (Management Ethernet [PTX Series Packet Transport Routers]) on page 1431
show interfaces brief (Management Ethernet) on page 1432
show interfaces brief (Management Ethernet [TX Matrix Plus Router]) on page 1432
show interfaces brief (Management Ethernet [PTX Series Packet Transport Routers]) on page 1432
show interfaces (Internal Ethernet) on page 1433
show interfaces (Internal Ethernet [TX Matrix Plus Router]) on page 1433
show interfaces detail (Internal Ethernet) on page 1434
show interfaces detail (Internal Ethernet [TX Matrix Plus Router]) on page 1434
show interfaces extensive (internal Ethernet) on page 1435
show interfaces extensive (internal Ethernet [TX Matrix Plus Router]) on page 1436

Output Fields Table 70 on page 1423 lists the output fields for the **show interfaces** (management) command on the M Series routers, T Series routers, TX Matrix Plus routers, and PTX Series. Output fields are listed in the approximate order in which they appear.

Table 70: show interfaces Output Fields for M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface

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 1110.	All levels
Interface index	Physical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Type	Type of interface.	All levels

Copyright © 2018, Juniper Networks, Inc.

Table 70: show interfaces Output Fields for M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface (continued)

Field Name	Field Description	Level of Output
Link-level type	Encapsulation type used on the physical interface.	All levels
MTU	Maximum transmission unit (MTU)—Size of the largest packet to be transmitted.	All levels
Clocking	Reference clock source of the interface.	All levels
Speed	Network speed on the interface.	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under “Common Output Fields Description” on page 1110.	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 1110.	All levels
Link type	Data transmission type.	**detail extensive** none
Link flags	Information about the link. Possible values are described in the “Link Flags” section under “Common Output Fields Description” on page 1110.	**detail extensive**
Physical info	Information about the physical interface.	**detail extensive**
Hold-times	Current interface hold-time up and hold-time down. Value is in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive** none
Hardware address	Media access control (MAC) address of the interface.	**detail extensive** none
Alternate link address	Backup link address.	**detail extensive**
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour*:*minute*:*second timezone*** **(*hour*:*minute*:*second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.	**detail extensive** none
Input packets	Number of packets received on the physical interface.	None specified
Output packets	Number of packets transmitted on the physical interface.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the logical and physical interface. • **Input bytes**, **Output bytes**—Number of bytes received and transmitted on the interface. • **Input packets**, **Output packets**—Number of packets received and transmitted on the interface.	**detail extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 70: show interfaces Output Fields for M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface (continued)

Field Name	Field Description	Level of Output
Input errors	• **Errors**—Input errors on the interface. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Frames received smaller than the runt threshold. • **Giants**—Frames received larger than the giant threshold. • **Policed Discards**—Frames that the incoming packet match code discarded because they were not recognized or were not of interest. Usually, this field reports protocols that Junos does not support. • **Resource errors**—Sum of transmit drops.	**extensive**
Output errors	• **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly, possibly once every 10 seconds, the cable, the remote system, or the interface is malfunctioning. • **Errors**—Sum of outgoing frame aborts and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet dropped by the ASIC RED mechanism. • **Resource errors**—Sum of transmit drops.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface	All levels
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	Logical interface SNMP interface index number.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface; values are described in the “Device Flags” section under “Common Output Fields Description” on page 1110.	All levels
Encapsulation	Encapsulation on the logical interface.	**detail extensive** none
inet	IP address of the logical interface.	**brief**
Protocol	Protocol family configured on the logical interface (such as **iso** or **inet6**).	**detail extensive** none
MTU	MTU size on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 70: show interfaces Output Fields for M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface (continued)

Field Name	Field Description	Level of Output
Route table	Route table in which this address exists. For example, **Route table:0** refers to inet.0.	**detail extensive**
Flags	Information about the protocol family flags. Possible values are described in the “Family Flags” section under “Common Output Fields Description” on page 1110.	**detail extensive** none
Addresses, Flags	Information about address flags. Possible values are described in the “Addresses Flags” section under “Common Output Fields Description” on page 1110.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces brief (Management Ethernet)

```
user@host> show interfaces fxp0 brief
Physical interface: fxp0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface fxp0.0
    Flags: SNMP-Traps Encapsulation: ENET2
    inet  192.168.70.143/21
```

show interfaces (Management Ethernet)

```
user@host> show interfaces fxp0
Physical interface: fxp0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 1
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Half-Duplex
  Current address: 00:00:5E:00:53:89, Hardware address: 00:00:5E:00:53:89
  Last flapped   : Never
    Input packets : 80804
    Output packets: 1105

  Logical interface fxp0.0 (Index 2) (SNMP ifIndex 13)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500
      Flags: Is-Primary
```

 Copyright © 2018, Juniper Networks, Inc.

```
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.64/21, Local: 192.168.70.143,
        Broadcast: 192.168.71.255
```

show interfaces (Management Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces em0
Physical interface: em0, Enabled, Physical link is Up
  Interface index: 8, SNMP ifIndex: 17
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Current address: 00:00:5E:00:53:c0, Hardware address: 00:00:5E:00:53:c0
  Last flapped   : Never
    Input packets : 1424
    Output packets: 5282

  Logical interface em0.0 (Index 3) (SNMP ifIndex 18)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 1424
    Output packets: 5282
    Protocol inet, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.178.0/25, Local: 192.168.178.11, Broadcast:
192.168.178.127
```

show interfaces (Management Ethernet [PTX Series Packet Transport Routers])

```
user@host> show interfaces em0
Physical interface: em0, Enabled, Physical link is Up
      Interface index: 8, SNMP ifIndex: 0
      Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps
      Device flags   : Present Running
      Interface flags: SNMP-Traps
      Link type      : Full-Duplex
      Current address: 00:00:5E:00:53:1b, Hardware address: 00:00:5E:00:53:1b
      Last flapped   : Never
 Input packets : 212581
 Output packets: 71

      Logical interface em0.0 (Index 3) (SNMP ifIndex 0)
 Flags: SNMP-Traps Encapsulation: ENET2
 Input packets : 212551
 Output packets: 71
 Protocol inet, MTU: 1500
   Flags: Is-Primary
   Addresses, Flags: Is-Default Is-Preferred Is-Primary
     Destination: 192.168.3/24, Local: 192.168.3.30,
     Broadcast: 192.168.3.255
```

show interfaces detail (Management Ethernet)

```
user@host> show interfaces fxp0 detail
Physical interface: fxp0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 1, Generation: 0
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
```

```
  Interface flags: SNMP-Traps
  Link type      : Half-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:89, Hardware address: 00:00:5E:00:53:89
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              6484031
   Output bytes  :               167503
   Input  packets:                81008
   Output packets:                 1110

  Logical interface fxp0.0 (Index 2) (SNMP ifIndex 13) (Generation 1)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Generation: 6, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.64/21, Local: 192.168.70.143,
        Broadcast: 192.168.71.255, Generation: 1
```

show interfaces detail (Management Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces em0 detail
Physical interface: em0, Enabled, Physical link is Up
  Interface index: 8, SNMP ifIndex: 17, Generation: 2
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:c0, Hardware address: 00:00:5E:00:53:c0
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :               124351
   Output bytes  :              1353212
   Input  packets:                 1804
   Output packets:                 5344
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0

  Logical interface em0.0 (Index 3) (SNMP ifIndex 18) (Generation 1)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :               117135
     Output bytes  :              1331647
     Input  packets:                 1804
     Output packets:                 5344
    Local statistics:
     Input  bytes  :               117135
     Output bytes  :              1331647
     Input  packets:                 1804
     Output packets:                 5344
```

 Copyright © 2018, Juniper Networks, Inc.

```
    Protocol inet, MTU: 1500, Generation: 1, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.178.0/25, Local: 192.168.178.11, Broadcast:
192.168.178.127, Generation: 1
```

show interfaces detail (Management Ethernet [PTX Packet Transport Routers])

```
user@host> show interfaces detail em0
Physical interface: em0, Enabled, Physical link is Up
      Interface index: 8, SNMP ifIndex: 0, Generation: 3
     Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,

      Speed: 1000mbps
      Device flags   : Present Running
      Interface flags: SNMP-Traps
      Link type      : Full-Duplex
      Physical info  : Unspecified
      Hold-times     : Up 0 ms, Down 0 ms
      Current address: 00:00:5E:00:53:1b, Hardware address: 00:00:5E:00:53:1b
      Alternate link address: Unspecified
      Last flapped   : Never
      Statistics last cleared: Never
      Traffic statistics:
       Input  bytes  :           15255909
       Output bytes  :               4608
       Input  packets:             214753
       Output packets:                 72
       IPv6 transit statistics:
 Input  bytes  :                  0
 Output bytes  :                  0
 Input  packets:                  0
 Output packets:                  0

      Logical interface em0.0 (Index 3) (SNMP ifIndex 0) (Generation 1)
 Flags: SNMP-Traps Encapsulation: ENET2
 Traffic statistics:
  Input  bytes  :           14394630
  Output bytes  :               3024
  Input  packets:             214723
  Output packets:                 72
 Local statistics:
  Input  bytes  :           14394630
  Output bytes  :               3024
  Input  packets:             214723
  Output packets:                 72
 Protocol inet, MTU: 1500, Generation: 1, Route table: 0
   Flags: Is-Primary
   Addresses, Flags: Is-Default Is-Preferred Is-Primary
     Destination: 192.168.3/24, Local: 192.168.3.30,
     Broadcast: 192.168.3.255, Generation: 1
```

show interfaces extensive (Management Ethernet)

```
user@host> show interfaces fxp0 extensive
Physical interface: fxp0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 1, Generation: 0
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
```

Copyright © 2018, Juniper Networks, Inc.

```
  Interface flags: SNMP-Traps
  Link type      : Half-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:89, Hardware address: 00:00:5E:00:53:89
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              6678904
   Output bytes  :               169657
   Input  packets:                83946
   Output packets:                 1127
  Input errors:
    Errors: 12, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0

  Logical interface fxp0.0 (Index 2) (SNMP ifIndex 13) (Generation 1)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Generation: 6, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.64/21, Local: 192.168.70.143,
        Broadcast: 192.168.71.255, Generation: 1
```

show interfaces extensive (Management Ethernet [TX Matrix Plus Router])

user@host> **show interfaces em0 extensive**

```
Physical interface: em0, Enabled, Physical link is Up
  Interface index: 8, SNMP ifIndex: 17, Generation: 2
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:c0, Hardware address: 00:00:5E:00:53:c0
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :               127120
   Output bytes  :              1357414
   Input  packets:                 1843
   Output packets:                 5372
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
   Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
 0
```

 Copyright © 2018, Juniper Networks, Inc.

```
  Logical interface em0.0 (Index 3) (SNMP ifIndex 18) (Generation 1)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :             119748
     Output bytes  :            1335719
     Input  packets:               1843
     Output packets:               5372
    Local statistics:
     Input  bytes  :             119748
     Output bytes  :            1335719
     Input  packets:               1843
     Output packets:               5372
    Protocol inet, MTU: 1500, Generation: 1, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.178.0/25, Local: 192.168.178.11, Broadcast:
192.168.178.127, Generation: 1
```

show interfaces extensive (Management Ethernet [PTX Series Packet Transport Routers])

```
user@host> show interfaces extensive em0
Physical interface: em0, Enabled, Physical link is Up
      Interface index: 8, SNMP ifIndex: 0, Generation: 3
     Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,

      Speed: 1000mbps
      Device flags   : Present Running
      Interface flags: SNMP-Traps
      Link type      : Full-Duplex
      Physical info  : Unspecified
      Hold-times     : Up 0 ms, Down 0 ms
      Current address: 00:00:5E:00:53:1b, Hardware address: 00:00:5E:00:53:1b
      Alternate link address: Unspecified
      Last flapped   : Never
      Statistics last cleared: Never
      Traffic statistics:
       Input  bytes  :           15236459
       Output bytes  :               4608
       Input  packets:             214482
       Output packets:                 72
       IPv6 transit statistics:
 Input  bytes  :                 0
 Output bytes  :                 0
 Input  packets:                 0
 Output packets:                 0
      Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
 Policed discards: 0, Resource errors: 0
      Output errors:
 Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
 Resource errors: 0

      Logical interface em0.0 (Index 3) (SNMP ifIndex 0) (Generation 1)
 Flags: SNMP-Traps Encapsulation: ENET2
 Traffic statistics:
  Input  bytes  :          14376264
  Output bytes  :              3024
  Input  packets:            214452
  Output packets:                72
 Local statistics:
```

Copyright © 2018, Juniper Networks, Inc.

```
 Input  bytes  :             14376264
 Output bytes  :                 3024
 Input  packets:               214452
 Output packets:                   72
Protocol inet, MTU: 1500, Generation: 1, Route table: 0
  Flags: Is-Primary
  Addresses, Flags: Is-Default Is-Preferred Is-Primary
    Destination: 192.168.3/24, Local: 192.168.3.30,
    Broadcast: 192.168.3.255, Generation: 1
```

show interfaces brief (Management Ethernet)

```
user@host> show interfaces fxp1 brief
Physical interface: fxp1, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface fxp1.0
    Flags: SNMP-Traps Encapsulation: ENET2
    inet  10.0.0.4/8
    inet6 fe80::200:ff:fe00:4/64
          fec0::10:0:0:4/64
    tnp   4
```

show interfaces brief (Management Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces em0 brief
Physical interface: em0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface em0.0
    Flags: SNMP-Traps Encapsulation: ENET2
    inet  192.168.178.11/25
```

show interfaces brief (Management Ethernet [PTX Series Packet Transport Routers])

```
user@host> show interfaces em0 brief
Physical interface: em0, Enabled, Physical link is Up
     Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,

      Speed: 1000mbps
      Device flags   : Present Running
      Interface flags: SNMP-Traps

      Logical interface em0.0
 Flags: SNMP-Traps Encapsulation: ENET2
 inet  192.168.3.30/24

    root@aboslutely> show interfaces em0 terse
    Interface               Admin Link Proto    Local                 Remote
    em0                     up    up
    em0.0                   up    up   inet     192.168.3.30/24
```

 Copyright © 2018, Juniper Networks, Inc.

show interfaces (Internal Ethernet)

```
user@host> show interfaces fxp1
Physical interface: fxp1, Enabled, Physical link is Up
  Interface index: 2, SNMP ifIndex: 2
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
  Last flapped   : Never
    Input packets : 30655
    Output packets: 33323

  Logical interface fxp1.0 (Index 3) (SNMP ifIndex 14)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 10/8, Local: 10.0.0.4, Broadcast: 10.255.255.255
    Protocol inet6, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::200:ff:fe00:4
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: fec0::/64, Local: fec0::10:0:0:4
    Protocol tnp, MTU: 1500
      Flags: Primary, Is-Primary
      Addresses
        Local: 4
```

show interfaces (Internal Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces ixgbe0
Physical interface: ixgbe0, Enabled, Physical link is Up
  Interface index: 2, SNMP ifIndex: 116
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
  Last flapped   : Never
    Input packets : 2301738
    Output packets: 3951155

  Logical interface ixgbe0.0 (Index 4) (SNMP ifIndex 117)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 2301595
    Output packets: 3951155
    Protocol inet, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: 10/8, Local: 10.34.0.4, Broadcast: 10.255.255.255
      Addresses, Flags: Primary Is-Default Is-Preferred Is-Primary
        Destination: 192.168/16, Local: 192.168.0.4, Broadcast: 192.168.0.4
    Protocol inet6, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::200:ff:fe22:4
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
```

Copyright © 2018, Juniper Networks, Inc.

```
        Destination: fec0::/64, Local: fec0::a:22:0:4
    Protocol tnp, MTU: 1500
      Flags: Primary, Is-Primary
      Addresses
        Local: 0x22000004
```

show interfaces detail (Internal Ethernet)

```
user@host> show interfaces fxp1 detail
Physical interface: fxp1, Enabled, Physical link is Up
  Interface index: 2, SNMP ifIndex: 2, Generation: 1
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              2339969
   Output bytes  :             15880707
   Input  packets:                30758
   Output packets:                33443

  Logical interface fxp1.0 (Index 3) (SNMP ifIndex 14) (Generation 2)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Generation: 7, Route table: 1
      Flags: Is-Primary
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 10/8, Local: 10.0.0.4, Broadcast: 10.255.255.255,
        Generation: 3
    Protocol inet6, MTU: 1500, Generation: 8, Route table: 1
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::200:ff:fe00:4,
        Broadcast: Unspecified, Generation: 5
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: fec0::/64, Local: fec0::10:0:0:4, Broadcast: Unspecified,
        Generation: 7
    Protocol tnp, MTU: 1500, Generation: 9, Route table: 1
      Flags: Primary, Is-Primary
      Addresses, Flags: None
        Destination: Unspecified, Local: 4, Broadcast: Unspecified,
        Generation: 8
```

show interfaces detail (Internal Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces ixgbe0 detail
Physical interface: ixgbe0, Enabled, Physical link is Up
  Interface index: 2, SNMP ifIndex: 116, Generation: 3
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
```

 Copyright © 2018, Juniper Networks, Inc.

```
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :            238172825
   Output bytes  :           1338948955
   Input  packets:              2360984
   Output packets:              4061512
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0

  Logical interface ixgbe0.0 (Index 4) (SNMP ifIndex 117) (Generation 2)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :            228720309
     Output bytes  :           1261387447
     Input  packets:              2360841
     Output packets:              4061512
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Local statistics:
     Input  bytes  :            228720309
     Output bytes  :           1261387447
     Input  packets:              2360841
     Output packets:              4061512
    Protocol inet, MTU: 1500, Generation: 2, Route table: 1
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
       Destination: 10/8, Local: 10.34.0.4, Broadcast: 10.255.255.255, Generation:
 2
      Addresses, Flags: Primary Is-Default Is-Preferred Is-Primary
        Destination: 192.168/16, Local: 192.168.0.4, Broadcast: 191.255.255.255,
 Generation: 3
    Protocol inet6, MTU: 1500, Generation: 3, Route table: 1
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::200:ff:fe22:4
    Generation: 4
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: fec0::/64, Local: fec0::a:22:0:4
    Protocol tnp, MTU: 1500, Generation: 5
    Generation: 4, Route table: 1
      Flags: Primary, Is-Primary
      Addresses, Flags: None
        Destination: Unspecified, Local: 0x22000004, Broadcast: Unspecified,
Generation: 6
```

show interfaces extensive (internal Ethernet)

```
user@host> show interfaces fxp1 extensive
Physical interface: fxp1, Enabled, Physical link is Up
  Interface index: 2, SNMP ifIndex: 2, Generation: 1
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
```

```
  Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              2349897
   Output bytes  :             15888605
   Input  packets:                30896
   Output packets:                33607
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0

  Logical interface fxp1.0 (Index 3) (SNMP ifIndex 14) (Generation 2)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Generation: 7, Route table: 1
      Flags: Is-Primary
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 10/8, Local: 10.0.0.4, Broadcast: 10.255.255.255,
        Generation: 3
    Protocol inet6, MTU: 1500, Generation: 8, Route table: 1
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::200:ff:fe00:4,
        Broadcast: Unspecified, Generation: 5
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: fec0::/64, Local: fec0::10:0:0:4, Broadcast: Unspecified,
        Generation: 7
    Protocol tnp, MTU: 1500, Generation: 9, Route table: 1
      Flags: Primary, Is-Primary
      Addresses, Flags: None
        Destination: Unspecified, Local: 4, Broadcast: Unspecified,
        Generation: 8
```

show interfaces extensive (internal Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces ixgbe0 extensive
Physical interface: ixgbe0, Enabled, Physical link is Up
  Interface index: 2, SNMP ifIndex: 116, Generation: 3
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :            242730780
```

 Copyright © 2018, Juniper Networks, Inc.

```
    Output bytes  :          1348312269
    Input  packets:             2398737
    Output packets:             4133510
    IPv6 transit statistics:
     Input  bytes  :                  0
     Output bytes  :                  0
     Input  packets:                  0
     Output packets:                  0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
 0

  Logical interface ixgbe0.0 (Index 4) (SNMP ifIndex 117) (Generation 2)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :           233127252
     Output bytes  :          1269350897
     Input  packets:             2398594
     Output packets:             4133510
     IPv6 transit statistics:
      Input  bytes  :                  0
      Output bytes  :                  0
      Input  packets:                  0
      Output packets:                  0
    Local statistics:
     Input  bytes  :           233127252
     Output bytes  :          1269350897
     Input  packets:             2398594
     Output packets:             4133510
    Protocol inet, MTU: 1500, Generation: 2, Route table: 1
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
       Destination: 10/8, Local: 10.34.0.4, Broadcast: 10.255.255.255, Generation:
 2
      Addresses, Flags: Primary Is-Default Is-Preferred Is-Primary
        Destination: 192.168/16, Local: 192.168.0.4, Broadcast: 191.255.255.255,
 Generation: 3
    Protocol inet6, MTU: 1500, Generation: 3, Route table: 1
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::200:ff:fe22:4
    Generation: 4
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: fec0::/64, Local: fec0::a:22:0:4
    Protocol tnp, MTU: 1500, Generation: 5
    Generation: 4, Route table: 1
      Flags: Primary, Is-Primary
      Addresses, Flags: None
        Destination: Unspecified, Local: 0x22000004, Broadcast: Unspecified,
Generation: 6
```

Copyright © 2018, Juniper Networks, Inc.

show interfaces (PPPoE)

Syntax

show interfaces pp0.*logical*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced before Junos OS Release 7.4.

Description

(M120 routers, M320 routers, and MX Series routers only). Display status information about the PPPoE interface.

Options

pp0.*logical*—Display standard status information about the PPPoE interface.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about PPPoE interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display PPPoE interface statistics.

Required Privilege Level

view

List of Sample Output

show interfaces (PPPoE) on page 1444
show interfaces (PPPoE over Aggregated Ethernet) on page 1444
show interfaces brief (PPPoE) on page 1445
show interfaces detail (PPPoE) on page 1445
show interfaces extensive (PPPoE on M120 and M320 Routers) on page 1446

Output Fields

Table 71 on page 1438 lists the output fields for the **show interfaces** (PPPoE) command. Output fields are listed in the approximate order in which they appear.

Table 71: show interfaces (PPPoE) Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the "Enabled Field" section under "Common Output Fields Description" on page 1110.	All levels
Interface index	Physical interface index number, which reflects its initialization sequence.	**detail extensive** none

Copyright © 2018, Juniper Networks, Inc.

Table 71: show interfaces (PPPoE) Output Fields (continued)

Field Name	Field Description	Level of Output
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Type	Physical interface type (PPPoE).	All levels
Link-level type	Encapsulation on the physical interface (PPPoE).	All levels
MTU	MTU size on the physical interface.	All levels
Clocking	Reference clock source. It can be **Internal** or **External**.	All levels
Speed	Speed at which the interface is running.	All levels
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under "Common Output Fields Description" on page 1110.	All levels
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under "Common Output Fields Description" on page 1110.	All levels
Link type	Physical interface link type: **full duplex** or **half duplex**.	All levels
Link flags	Information about the interface. Possible values are described in the "Link Flags" section under "Common Output Fields Description" on page 1110.	All levels
Input rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output rate	Output rate in bps and pps.	None specified
Physical Info	Physical interface information.	All levels
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive**
Hardware address	MAC address of the hardware.	**detail extensive**
Alternate link address	Backup address of the link.	**detail extensive**
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**

Table 71: show interfaces (PPPoE) Output Fields (continued)

Field Name	Field Description	Level of Output
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled. NOTE: These fields include dropped traffic and exception traffic, as those fields are not separately defined. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Input errors	Input errors on the interface: • **Errors**—Sum of incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Giants**—Number of frames received that are larger than the giant threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • **Resource errors**—Sum of B chip Tx drops and IXP Tx net transmit drops.	**extensive**
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions** —Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), then the cable, the far-end system, or the PIM is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of B chip Tx drops and IXP Tx net transmit drops.	**extensive**
Logical Interface		

 Copyright © 2018, Juniper Networks, Inc.

Table 71: show interfaces (PPPoE) Output Fields (continued)

Field Name	Field Description	Level of Output
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number (which reflects its initialization sequence).	**detail extensive** none
SNMP ifIndex	Logical interface SNMP interface index number.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under "Common Output Fields Description" on page 1110.	All levels
Encapsulation	Type of encapsulation configured on the logical interface.	All levels
PPP parameters	PPP status: • LCP restart timer—Length of time (in milliseconds) between successive Link Control Protocol (LCP) configuration requests. • NCP restart timer—Length of time (in milliseconds) between successive Network Control Protocol (NCP) configuration requests.	**detail**
PPPoE	PPPoE status: • **State**—State of the logical interface (**up** or **down**). • **Session ID**—PPPoE session ID. • **Service name**—Type of service required. Can be used to indicate an Internet service provider (ISP) name or a class or quality of service. • **Configured AC name**—Configured access concentrator name. • **Auto-reconnect timeout**—Time after which to try to reconnect after a PPPoE session is terminated, in seconds. • **Idle Timeout**—Length of time (in seconds) that a connection can be idle before disconnecting. • **Underlying interface**—Interface on which PPPoE is running.	All levels
Link	Name of the physical interfaces for member links in an aggregated Ethernet bundle for a PPPoE over aggregated Ethernet configuration. PPPoE traffic goes out on these interfaces.	All levels
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. This counter usually takes less than 1 second to stabilize.	**detail extensive**

Table 71: show interfaces (PPPoE) Output Fields (continued)

Field Name	Field Description	Level of Output
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled. NOTE: The packet and byte counts in these fields include traffic that is dropped and does not leave the router. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. This counter usually takes less than 1 second to stabilize.	**detail extensive**
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. This counter usually takes less than 1 second to stabilize. NOTE: The packet and byte counts in these fields include traffic that is dropped and does not leave the router.	**detail extensive**
Keepalive settings	(PPP and HDLC) Configured settings for keepalives. • **interval** ***seconds***—The time in seconds between successive keepalive requests. The range is **10** seconds through **32,767** seconds, with a default of **10** seconds. • **down-count*number***—The number of keepalive packets a destination must fail to receive before the network takes a link down. The range is **1** through **255**, with a default of **3**. • **up-count** ***number***—The number of keepalive packets a destination must receive to change a link's status from down to up. The range is **1** through **255**, with a default of **1**.	**detail extensive**

Keepalive statistics	(PPP and HDLC) Information about keepalive packets.
	
	• **Input**—Number of keepalive packets received by PPP.
	  ◦ **(last seen 00:00:00 ago)**—Time the last keepalive packet was received, in the format ***hh*:*mm*:*ss***.
	• **Output**—Number of keepalive packets sent by PPP and how long ago the last keepalive packets were sent and received.
	  ◦ **(last seen 00:00:00 ago)**—Time the last keepalive packet was sent, in the format ***hh*:*mm*:*ss***.
	
	(MX Series routers with MPCs/MICs) When an MX Series router with MPCs/MICs is using PPP fast keepalive for a PPP link, the display does not include the number of keepalive packets received or sent, or the amount of time since the router received or sent the last keepalive packet.
		detail extensive

Input packets	Number of packets received on the logical interface.	None specified
Output packets	Number of packets transmitted on the logical interface.	None specified

 Copyright © 2018, Juniper Networks, Inc.

Table 71: show interfaces (PPPoE) Output Fields (continued)

Field Name	Field Description	Level of Output
LCP state	(PPP) Link Control Protocol state. • **Conf-ack-received**—Acknowledgement was received. • **Conf-ack-sent**—Acknowledgement was sent. • **Conf-req-sent**—Request was sent. • **Down**—LCP negotiation is incomplete (not yet completed or has failed). • **Not-configured**—LCP is not configured on the interface. • **Opened**—LCP negotiation is successful.	none **detail extensive**
NCP state	(PPP) Network Control Protocol state. • **Conf-ack-received**—Acknowledgement was received. • **Conf-ack-sent**—Acknowledgement was sent. • **Conf-req-sent**—Request was sent. • **Down**—NCP negotiation is incomplete (not yet completed or has failed). • **Not-configured**—NCP is not configured on the interface. • **Opened**—NCP negotiation is successful.	**detail extensive** none
CHAP state	(PPP) Displays the state of the Challenge Handshake Authentication Protocol (CHAP) during its transaction. • **Chap-Chal-received**—Challenge was received but response not yet sent. • **Chap-Chal-sent**—Challenge was sent. • **Chap-Resp-received**—Response was received for the challenge sent, but CHAP has not yet moved into the Success state. (Most likely with RADIUS authentication.) • **Chap-Resp-sent**—Response was sent for the challenge received. • **Closed**—CHAP authentication is incomplete. • **Failure**—CHAP authentication failed. • **Not-configured**—CHAP is not configured on the interface. • **Success**—CHAP authentication was successful.	none **detail extensive**
Protocol	Protocol family configured on the logical interface.	**detail extensive** none
protocol-family	Protocol family configured on the logical interface. If the protocol is **inet**, the IP address of the interface is also displayed.	**brief**
MTU	MTU size on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route table	Routing table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive** none
Flags	Information about the protocol family flags. Possible values are described in the "Family Flags" section under "Common Output Fields Description" on page 1110.	**detail extensive** none

Copyright © 2018, Juniper Networks, Inc.

Table 71: show interfaces (PPPoE) Output Fields (continued)

Field Name	Field Description	Level of Output
Addresses, Flags	Information about the addresses configured for the protocol family. Possible values are described in the “Addresses Flags” section under “Common Output Fields Description” on page 1110.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address.	**detail extensive** none

Sample Output

show interfaces (PPPoE)

```
user@host> show interfaces pp0
Physical interface: pp0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 24
  Type: PPPoE, Link-level type: PPPoE, MTU: 1532
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface pp0.0 (Index 72) (SNMP ifIndex 72)
    Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
    PPPoE:
      State: SessionDown, Session ID: None,
      Service name: None, Configured AC name: sapphire,
      Auto-reconnect timeout: 100 seconds, Idle timeout: Never,
      Underlying interface: at-5/0/0.0 (Index 70)
  Input packets : 0
  Output packets: 0
  LCP state: Not-configured
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
  mpls: Not-configured
  CHAP state: Closed
    Protocol inet, MTU: 100
      Flags: User-MTU, Negotiate-Address
```

show interfaces (PPPoE over Aggregated Ethernet)

```
user@host> show interfaces pp0.1073773821
Logical interface pp0.1073773821 (Index 80) (SNMP ifIndex 32584)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
    PPPoE:
      State: SessionUp, Session ID: 1,
      Session AC name: alcor, Remote MAC address: 00:00:5e:00:53:01,
      Underlying interface: demux0.100 (Index 88)
    Link:
      ge-1/0/0.32767
      ge-1/0/1.32767
```


Copyright © 2018, Juniper Networks, Inc.

```
    Input packets : 6
    Output packets: 6
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured,  mpls:
Not-configured
  CHAP state: Closed
  PAP state: Success
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Primary
        Local: 203.0.113.1
```

show interfaces brief (PPPoE)

```
user@host> show interfaces pp0 brief
Physical interface: pp0, Enabled, Physical link is Up
  Type: PPPoE, Link-level type: PPPoE, MTU: 1532, Speed: Unspecified
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps

  Logical interface pp0.0
    Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
    PPPoE:
      State: SessionDown, Session ID: None,
      Service name: None, Configured AC name: sapphire,
      Auto-reconnect timeout: 100 seconds, Idle timeout: Never,
      Underlying interface: at-5/0/0.0 (Index 70)
    inet
```

show interfaces detail (PPPoE)

```
user@host> show interfaces pp0 detail
Physical interface: pp0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 24, Generation: 9
  Type: PPPoE, Link-level type: PPPoE, MTU: 1532, Speed: Unspecified
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
  Logical interface pp0.0 (Index 72) (SNMP ifIndex 72) (Generation 14)
    Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
    PPPoE:
      State: SessionDown, Session ID: None,
      Service name: None, Configured AC name: sapphire,
      Auto-reconnect timeout: 100 seconds, Idle timeout: Never,
      Underlying interface: at-5/0/0.0 (Index 70)
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
```

```
      Output packets:                   0
     Local statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
     Transit statistics:
      Input  bytes  :                   0                    0 bps
      Output bytes  :                   0                    0 bps
      Input  packets:                   0                    0 pps
      Output packets:                   0                    0 pps
   LCP state: Not-configured
   NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
   mpls: Not-configured
   CHAP state: Closed
     Protocol inet, MTU: 100, Generation: 14, Route table: 0
       Flags: User-MTU, Negotiate-Address
```

show interfaces extensive (PPPoE on M120 and M320 Routers)

```
user@host> show interfaces pp0 extensive
Physical interface: pp0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 93, Generation: 129
  Type: PPPoE, Link-level type: PPPoE, MTU: 1532, Speed: Unspecified
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :               972192                    0 bps
   Output bytes  :               975010                    0 bps
   Input  packets:                 1338                    0 pps
   Output packets:                 1473                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
 0,
    Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
 0

  Logical interface pp0.0 (Index 69) (SNMP ifIndex 96) (Generation 194)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
    PPPoE:
      State: SessionUp, Session ID: 26,
      Session AC name: None, AC MAC address: 00:00:5e:00:53:12,
      Service name: None, Configured AC name: None,
      Auto-reconnect timeout: Never, Idle timeout: Never,
      Underlying interface: ge-3/0/1.0 (Index 67)
    Traffic statistics:
     Input  bytes  :                  252
```

 Copyright © 2018, Juniper Networks, Inc.

```
     Output bytes  :                296
     Input  packets:                  7
     Output packets:                  8
     IPv6 transit statistics:
      Input  bytes  :                 0
      Output bytes  :                 0
      Input  packets:                 0
      Output packets:                 0
    Local statistics:
     Input  bytes  :                252
     Output bytes  :                296
     Input  packets:                  7
     Output packets:                  8
    Transit statistics:
     Input  bytes  :                  0                    0 bps
     Output bytes  :                  0                    0 bps
     Input  packets:                  0                    0 pps
     Output packets:                  0                    0 pps
     IPv6 transit statistics:
      Input  bytes  :                 0
      Output bytes  :                 0
      Input  packets:                 0
      Output packets:                 0
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive statistics:
    Input : 1 (last seen 00:00:00 ago)
    Output: 1 (last sent 00:00:03 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
  CHAP state: Closed
  PAP state: Closed
    Protocol inet, MTU: 1492, Generation: 171, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 203.0.113.2, Local: 203.0.113.1, Broadcast: Unspecified,
Generation: 206
```

show interfaces (PTX Series Packet Transport Routers)

Syntax

show interfaces et-*fpc/pic/port*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced in Junos OS Release 8.0.
Command introduced in Junos OS Release 12.1 for PTX Series Packet Transport Routers.

Description

(PTX Series Packet Transport Routers only) Display status information about the specified Ethernet interface.

Options

et-*fpc/pic/port*—Display standard information about the specified Ethernet interface.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media —(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

List of Sample Output

show interfaces brief (PTX5000 Packet Transport Router) on page 1458
show interfaces extensive (PTX5000 Packet Transport Router) on page 1458
show interfaces terse (PTX5000 Packet Transport Router) on page 1459

Output Fields

See Table 72 on page 1449 for the output fields for the **show interfaces** (PTX Series Packet Transport Routers) command.

Copyright © 2018, Juniper Networks, Inc.

Table 72: show interfaces PTX Series 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 1110.	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
BPDU Error	Bridge protocol data unit (BPDU) errors (if any).	All levels
MAC-Rewrite	MAC Rewrite errors (if any).	All levels
Loopback	Loopback status: **Enabled** or **Disabled**. If loopback is enabled, type of loopback: **Local** or **Remote**.	All levels
Source filtering	Source filtering status: **Enabled** or **Disabled**.	All levels
Flow control	Flow control status: **Enabled** or **Disabled**.	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under “Common Output Fields Description” on page 1110.	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 1110.	All levels
Link flags	Information about the link. Possible values are described in the “Links Flags” section under “Common Output Fields Description” on page 1110.	All levels
CoS queues	Number of CoS queues configured.	**detail extensive** none
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

Copyright © 2018, Juniper Networks, Inc.

Table 72: show interfaces PTX Series Output Fields (continued)

Field Name	Field Description	Level of Output
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour:minute:second:timezone*** **(*hour:minute:second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).**	**detail extensive** none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface. NOTE: **Input bytes** and **output bytes** are counted as Layer 3 packet length.	**detail extensive**
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • **L3 incompletes**—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the **ignore-l3-incompletes** statement. NOTE: The L3 incompletes field is *not* supported on PTX Series Packet Transport Routers. • **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **FIFO errors**—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **Resource errors**—Sum of transmit drops.	**extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 72: show interfaces PTX Series Output Fields (continued)

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Collisions**—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • **Aged packets**—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • **FIFO errors**—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**
Egress queues	Total number of egress queues supported on the specified interface.	**detail extensive**
Queue counters (Egress)	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**
Ingress queues	Total number of ingress queues supported on the specified interface.	**extensive**
Queue counters (Ingress)	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 72: show interfaces PTX Series Output Fields (continued)

Field Name	Field Description	Level of Output
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value **None** or **Link**. • **None**—There are no active defects or alarms. • **Link**—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. • **LOCAL-FAULT**—Link fault signaling operates between the remote PHY RS (Reconciliation sub-layer) and the local RS. A Local Fault is used to signal a detected fault between the remote RS and the local RS to the local Ethernet interface. • **REMOTE-FAULT**—When the Local Fault status reaches an RS, the RS stops sending MAC data and continuously generates the Remote Fault status on the transmit data path .	**detail extensive** none
MAC statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC subsystem, including the following: • **Total octets** and **total packets**—Total number of octets and packets. • **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**

 Copyright © 2018, Juniper Networks, Inc.

Table 72: show interfaces PTX Series Output Fields (continued)

Field Name	Field Description	Level of Output
Filter statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected. • **Input packet count**—Number of packets received from the MAC hardware that the filter processed. • **Input packet rejects**—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • **Input DA rejects**—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the router from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local router (which the router is rejecting). • **Input SA rejects**—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • **Output packet count**—Number of packets that the filter has given to the MAC hardware. • **Output packet pad count**—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • **Output packet error count**—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment. • **CAM destination filters, CAM source filters**—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields should be 0.	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 72: show interfaces PTX Series Output Fields (continued)

Field Name	Field Description	Level of Output

Autonegotiation information	Information about link autonegotiation.
	• **Negotiation status:**
	  ◦ **Incomplete**—Ethernet interface has the speed or link mode configured.
	  ◦ **No autonegotiation**—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
	  ◦ **Complete**—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• **Link partner status**—**OK** when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• **Link partner:**
	  ◦ **Link mode**—Depending on the capability of the attached Ethernet device, either **Full-duplex** or **Half-duplex**.
	  ◦ **Flow control**—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is **None**. For Gigabit Ethernet interfaces, types are **Symmetric** (link partner supports **PAUSE** on receive and transmit), **Asymmetric** (link partner supports **PAUSE** on transmit), and **Symmetric/Asymmetric** (link partner supports both **PAUSE** on receive and transmit or only **PAUSE** receive).
	  ◦ **Remote fault**—Remote fault information from the link partner—**Failure** indicates a receive link error. **OK** indicates that the link partner is receiving. **Negotiation error** indicates a negotiation error. **Offline** indicates that the link partner is going offline.
	• **Local resolution**—Information from the link partner:
	  ◦ **Flow control**—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are **Symmetric** (link partner supports **PAUSE** on receive and transmit), **Asymmetric** (link partner supports **PAUSE** on transmit), and **Symmetric/Asymmetric** (link partner supports both **PAUSE** on receive and transmit or only **PAUSE** receive).
	  ◦ **Remote fault**—Remote fault information. **Link OK** (no error detected on receive), **Offline** (local interface is offline), and **Link Failure** (link error detected on receive).
		extensive

Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number.	**extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 72: show interfaces PTX Series Output Fields (continued)

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface. • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP interface index number for the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under “Common Output Fields Description” on page 1110.	All levels

Table 72: show interfaces PTX Series Output Fields (continued)

Field Name	Field Description	Level of Output
VLAN-Tag	Rewrite profile applied to incoming or outgoing frames on the outer (**Out**) VLAN tag or for both the outer and inner (**In**) VLAN tags. • **push**—An outer VLAN tag is pushed in front of the existing VLAN tag. • **pop**—The outer VLAN tag of the incoming frame is removed. • **swap**—The outer VLAN tag of the incoming frame is overwritten with the user-specified VLAN tag information. • **push**—An outer VLAN tag is pushed in front of the existing VLAN tag. • **push-push**—Two VLAN tags are pushed in from the incoming frame. • **swap-push**—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • **swap-swap**—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user-specified VLAN tag value. • **pop-swap**—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • **pop-pop**—Both the outer and inner VLAN tags of the incoming frame are removed.	**brief detail extensive** none
Demux	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: • Source Family Inet • Destination Family Inet	**detail extensive** none
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under “Common Output Fields Description” on page 1110.	**detail extensive** none
MTU	Maximum transmission unit size on the logical interface.	**detail extensive** none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	**detail extensive** none
Traffic statistics	Number and rate of bytes and packets received and transmitted on the specified interface set. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface set • **Input packets, Output packets**—Number of packets received and transmitted on the interface set. NOTE: **Input bytes** and **output bytes** are counted as Layer 3 packet length.	**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**

 Copyright © 2018, Juniper Networks, Inc.

Table 72: show interfaces PTX Series Output Fields (continued)

Field Name	Field Description	Level of Output
Transit statistics	Number and rate of bytes and packets transiting the switch.	**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 1110.	**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**
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 1110.	**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 flags (possible values are described in the “Addresses Flags” section under “Common Output Fields Description” on page 1110.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interlace.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show interfaces brief (PTX5000 Packet Transport Router)

```
user@host> show interfaces brief et-7/0/0
Physical interface: et-7/0/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Loopback: Disabled, Source
 filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
```

show interfaces extensive (PTX5000 Packet Transport Router)

```
user@host> show interfaces et-7/0/0 extensive
Physical interface: et-7/0/0, Enabled, Physical link is Up
  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 (01:51:00 ago)
  Statistics last cleared: 2012-01-18 13:38:54 PST (00:00:30 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, 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
```

 Copyright © 2018, Juniper Networks, Inc.

```
    3                   network-control
  Active alarms  : None
  Active defects : None
  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
  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
    0 best-effort            95     9500000000    95               0       low
none
    3 network-control         5      500000000     5               0       low
none
  Interface transmit statistics: Disabled
```

show interfaces terse (PTX5000 Packet Transport Router)

```
user@host> show interfaces terse
Interface               Admin Link Proto    Local                 Remote
et-2/0/0                up    up
et-2/0/1                up    up
et-2/0/2                up    up
et-2/0/3                up    up
et-2/0/4                up    up
et-2/0/5                up    down
et-2/0/6                up    up
et-2/0/7                up    up
et-2/0/8                up    up
et-2/0/9                up    down
et-2/0/10               up    up
et-2/0/11               up    up
et-2/0/12               up    up
et-2/0/13               up    down
et-2/0/14               up    up
```

Copyright © 2018, Juniper Networks, Inc.

```
et-2/0/15               up    up
et-2/0/16               up    up
et-2/0/17               up    down
et-2/0/18               up    down
et-2/0/19               up    up
et-2/0/20               up    down
et-2/0/21               up    up
et-2/0/22               up    down
et-2/0/23               up    up
et-2/1/0                up    up
et-2/1/1                up    up
et-2/1/2                up    up
et-2/1/3                up    up
et-2/1/4                up    up
et-2/1/5                up    up
et-2/1/6                up    up
et-2/1/7                up    up
et-2/1/8                up    up
et-2/1/9                up    up
et-2/1/10               up    up
et-2/1/11               up    up
et-2/1/12               up    up
et-2/1/13               up    up
et-2/1/14               up    up
et-2/1/15               up    up
et-2/1/16               up    up
et-2/1/17               up    up
et-2/1/18               up    up
et-2/1/19               up    up
et-2/1/20               up    up
et-2/1/21               up    up
et-2/1/22               up    up
et-2/1/23               up    up
et-5/0/0                up    up
et-5/0/0.0              up    up   ccc
et-5/0/0.32767          up    up   multiservice
et-5/0/1                up    up
et-5/0/2                up    up
et-5/0/3                up    down
et-5/0/4                up    down
et-5/0/5                up    up
et-5/0/5.0              up    up   ccc
et-5/0/5.32767          up    up   multiservice
et-5/0/6                up    up
et-5/0/7                up    up
et-5/0/8                up    down
et-5/0/9                up    up
et-5/0/10               up    up
et-5/0/11               up    up
et-5/0/12               up    up
et-5/0/13               up    down
et-5/0/14               up    down
et-5/0/15               up    up
et-5/0/16               up    up
et-5/0/17               up    up
et-5/0/18               up    up
et-5/0/19               up    up
et-5/0/20               up    down
et-5/0/21               up    down
et-5/0/22               up    up
et-5/0/23               up    up
```

 Copyright © 2018, Juniper Networks, Inc.

```
et-5/1/0                up    up
et-5/1/1                up    up
et-7/0/0                up    up
et-7/0/1                up    up
et-7/0/2                up    up
et-7/0/3                up    up
et-7/0/4                up    up
et-7/0/5                up    up
et-7/0/6                up    up
et-7/0/7                up    up
et-7/0/8                up    up
et-7/0/9                up    up
et-7/0/10               up    down
et-7/0/11               up    down
et-7/0/12               up    down
et-7/0/13               up    down
et-7/0/14               up    down
et-7/0/15               up    down
et-7/0/16               up    down
et-7/0/17               up    down
et-7/0/18               up    down
et-7/0/19               up    down
et-7/0/20               up    down
et-7/0/21               up    down
et-7/0/22               up    down
et-7/0/23               up    down
dsc                     up    up
em0                     up    up
em0.0                   up    up   inet     192.168.177.61/25
gre                     up    up
ipip                    up    up
ixgbe0                  up    up
ixgbe0.0                up    up   inet     10.0.0.4/8
                                            128.0.0.1/2
                                            128.0.0.4/2
                                   inet6    fe80::200:ff:fe00:4/64
                                            fec0::a:0:0:4/64
                                   tnp      0x4
ixgbe1                  up    up
ixgbe1.0                up    up   inet     10.0.0.4/8
                                            128.0.0.1/2
                                            128.0.0.4/2
                                   inet6    fe80::200:1ff:fe00:4/64
                                            fec0::a:0:0:4/64
                                   tnp      0x4
lo0                     up    up
lo0.0                   up    up   inet     10.255.177.61       --> 0/0
                                            127.0.0.1           --> 0/0
                                   iso
47.0005.80ff.f800.0000.0108.0001.0102.5517.7061
                                   inet6    abcd::10:255:177:61
                                            fe80::ee9e:cd0f:fc02:b01e
lo0.16384               up    up   inet     127.0.0.1           --> 0/0
lo0.16385               up    up   inet
lsi                     up    up
mtun                    up    up
pimd                    up    up
pime                    up    up
tap                     up    up
```

show interfaces (SONET/SDH)

Syntax

show interfaces so-*fpc/pic/port*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced before Junos OS Release 7.4.

Description

(M Series and T Series routers only) Display status information about the specified SONET/SDH interface.

Options

so-*fpc/pic/port*—Display standard information about the specified SONET/SDH interface.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Related Documentation

- *SONET/SDH Interfaces Overview*

List of Sample Output

show interfaces (SDH Mode, PPP) on page 1475
show interfaces brief (SDH Mode, PPP) on page 1476
show interfaces detail (SDH Mode, PPP) on page 1476
show interfaces extensive (SDH Mode, PPP) on page 1477
show interfaces brief (SONET Mode, Frame Relay) on page 1479
show interfaces (SONET Mode, Frame Relay) on page 1479
show interfaces detail (SONET Mode, Frame Relay) on page 1480
show interfaces extensive (SONET Mode, Frame Relay) on page 1482
show interfaces extensive (OC768-over-4xOC192 Mode) on page 1484
show interfaces detail (IPv6 Tracking) on page 1487
show interfaces (Shared Interface) on page 1488

Output Fields

Table 73 on page 1463 lists the output fields for the **show interfaces** (SONET/SDH) command. Output fields are listed in the approximate order in which they appear.

Copyright © 2018, Juniper Networks, Inc.

Table 73: SONET/SDH show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under “Common Output Fields Description” on page 1110.	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**
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	MTU size on the physical interface.	All levels
Clocking	SONET/SDH reference clock source: **Internal** or **External**. Clocking is configured and displayed only for channel 0.	All levels
***Framing* mode**	Framing mode: **SONET** or **SDH**.	All levels
Speed	Speed at which the interface is running.	All levels
Loopback	Whether loopback is enabled and the type of loopback (**local** or **remote**).	All levels
FCS	Frame check sequence on the interface (either **16** or **32**). The default is **16** bits.	All levels
Payload scrambler	Whether payload scrambling is enabled.	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 1110.	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 1110.	All levels
Shared-interface	Indicates whether the routing domain is the owner or non-owner of the shared interface. If the routing domain is the Root System Domain (RSD), the value is **Owner**. If the routing domain is a Protected System Domain (PSD) under the RSD, the value is **Non-owner**.	All levels
Link flags	Information about the link. Possible values are described in the “Link Flags” section under “Common Output Fields Description” on page 1110.	All levels

Table 73: SONET/SDH show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
ANSI or ITU LMI settings	(Frame Relay) Settings for Local Management Interface (LMI). The format is **(ANSI or ITU) LMI settings:** *value*, *value*... *xx* seconds, where *value* can be: • **n391dte**—DTE full status polling interval (1-255) • **n392dce**—DCE error threshold (1-10) • **n392dte**—DTE error threshold (1-10) • **n393dce**—DCE monitored event count (1-10) • **n393dte**—DTE monitored event count (1-10) • **t391dte**—DTE polling timer (5-30 seconds) • **t392dce**—DCE polling verification timer (5-30 seconds)	All levels
LMI	**Input:** ***value*** **(*hh*:*mm*:*ss* ago)**, **Output:** ***value*** **(*hh*:*mm*:*ss* ago)**	**brief** none
LMI statistics	(Frame Relay) LMI packet statistics: • **Input**—Number of packets coming in on the interface (*nn*) and how much time has passed since the last packet arrived. The format is **Input:** ***nn*** **(last seen *hh*:*mm*:*ss* ago)**. • **Output**—Number of packets sent out on the interface (*nn*) and how much time has passed since the last packet was sent. The format is **Output:** ***nn*** **(last sent *hh*:*mm*:*ss* ago)**.	**detail extensive**
DTE statistics	(Frame Relay) Statistics about messages transmitted from the data terminal equipment (DTE) to the data circuit-terminating equipment (DCE): • **Enquiries sent**—Number of link status enquiries sent from the DTE to the DCE. • **Full enquiries sent**—Number of full enquiries sent from the DTE to the DCE. • **Enquiry responses received**—Number of enquiry responses received by the DTE from the DCE. • **Full enquiry responses received**—Number of full enquiry responses sent from the DTE to the DCE.	**detail extensive** none
DCE statistics	(Frame Relay) Statistics about messages transmitted from the DCE to the DTE: • **Enquiries received**—Number of enquiries received by the DCE from the DTE. • **Full enquiries received**—Number of full enquiries received by the DCE from the DTE. • **Enquiry responses sent**—Number of enquiry responses sent from the DCE to the DTE. • **Full enquiry responses sent**—Number of full enquiry responses sent from the DCE to the DTE.	**detail extensive** none

 Copyright © 2018, Juniper Networks, Inc.

Table 73: SONET/SDH show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Common statistics	(Frame Relay) Statistics about messages sent between the DTE and the DCE: • **Unknown messages received**—Number of received packets that do not fall into any category. • **Asynchronous updates received**—Number of link status peer changes received. • **Out-of-sequence packets received**—Number of packets for which the sequence of the packets received is different from the expected sequence. • **Keepalive responses timedout**—Number of keepalive responses that timed out when no LMI packet was reported for **n392dte** or **n393dce** intervals. (See **LMI settings**.)	**detail extensive** none
Nonmatching DCE-end DLCIs	(Frame Relay. Displayed only from the DTE) Number of DLCIs configured from the DCE.	**detail extensive**
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Keepalive settings	(PPP and HDLC) Configured settings for keepalives. • **interval** ***seconds***—The time in seconds between successive keepalive requests. The range is **10** seconds through **32,767** seconds, with a default of **10** seconds. • **down-count** ***number***—The number of keepalive packets a destination must fail to receive before the network takes a link down. The range is **1** through **255**, with a default of **3**. • **up-count** ***number***—The number of keepalive packets a destination must receive to change a link's status from down to up. The range is **1** through **255**, with a default of **1**.	All levels

Keepalive or **Keepalive statistics**	(PPP and HDLC) Information about keepalive packets.
	• **Input**—Number of keepalive packets received by PPP.
	  ◦ **(last seen 00:00:00 ago)**—Time since the last keepalive packet was received, in the format ***hh:mm:ss***.
	• **Output**—Number of keepalive packets sent by PPP and how long ago the last keepalive packets were sent and received.
	  ◦ **(last seen 00:00:00 ago)**—Time since the last keepalive packet was sent, in the format ***hh:mm:ss***.
		All levels

LCP state	(PPP) Link Control Protocol state. • **Conf-ack-received**—Acknowledgement was received. • **Conf-ack-sent**—Acknowledgement was sent. • **Conf-req-sent**—Request was sent. • **Down**—LCP negotiation is incomplete (not yet completed or has failed). • **Not-configured**—LCP is not configured on the interface. • **Opened**—LCP negotiation is successful.	**detail extensive** none

Copyright © 2018, Juniper Networks, Inc.

Table 73: SONET/SDH show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
NCP state	(PPP) Network Control Protocol state. • **Conf-ack-received**—Acknowledgement was received. • **Conf-ack-sent**—Acknowledgement was sent. • **Conf-req-sent**—Request was sent. • **Down**—NCP negotiation is incomplete (not yet completed or has failed). • **Not-configured**—NCP is not configured on the interface. • **Opened**—NCP negotiation is successful.	**detail extensive** none
CHAP state	(PPP) Displays the state of the Challenge Handshake Authentication Protocol (CHAP) during its transaction. • **Chap-Chal-received**—Challenge was received but response not yet sent. • **Chap-Chal-sent**—Challenge was sent. • **Chap-Resp-received**—Response was received for the challenge sent, but CHAP has not yet moved into the **Success** state. (Most likely with RADIUS authentication.) • **Chap-Resp-sent**—Response was sent for the challenge received. • **Closed**—CHAP authentication is incomplete. • **Failure**—CHAP authentication failed. • **Not-configured**—CHAP is not configured on the interface. • **Success**—CHAP authentication was successful.	**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
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number of bytes and packets received and transmitted on the physical interface, and the traffic rate in bits per seconds (bps). • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 73: SONET/SDH show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Label-switched interface (LSI) traffic statistics	(Frame Relay) LSI traffic statistics: • **Input bytes**—Number of bytes and speed, in bits per second (bps), received on the interface. • **Output packets**—Number of packets and speed, in bps, transmitted on the interface.	**extensive**
Input errors	Input errors on the interface whose definitions are as follows: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Giants**—Number of frames received that are larger than the giant threshold. • **Bucket Drops**—Drops resulting from the traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is **off**. • **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. • **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. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • **HS link FIFO overflows**—Number of FIFO overflows on the high-speed links between the ASICs responsible for handling the router interfaces.	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 73: SONET/SDH show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and 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. • **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. • **HS link FIFO underflows**—Number of FIFO underflows on the high-speed links between the ASICs responsible for handling the router interfaces. • **MTU errors**—Number of packets whose size exceeds the MTU of the interface.	**extensive**
IPv6 transit statistics	Number of transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Egress queues	Total number of egress queues supported on the specified interface.	**detail extensive**
Queue counters	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**extensive**
SONET alarms **SONET defects**	(SONET) SONET media-specific alarms and defects that prevents the interface from passing packets. When a defect persists for a certain period, 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 light the red or yellow alarm LED on the craft interface. See these fields for possible alarms and defects: **SONET PHY**, **SONET section**, **SONET line**, and **SONET path**.	All levels
Link	(For 4-port OC192c PIC operating in OC768-over-4xOC192 mode) The link number. Errors and alarms are displayed for each link.	**extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 73: SONET/SDH show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
SONET PHY	Counts of specific SONET errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **PLL Lock**—Phase-locked loop • **PHY Light**—Loss of optical signal	**extensive**
SONET section	Counts of specific SONET errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **BIP-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**
SONET line	Active alarms and defects, plus counts of specific SONET errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **BIP-B2**—Bit interleaved parity for SONET line overhead • **REI-L**—Remote error indication (near-end line) • **RDI-L**—Remote defect indication (near-end line) • **AIS-L**—Alarm indication signal (near-end line) • **BERR-SF**—Bit error rate fault (signal failure) • **BERR-SD**—Bit error rate defect (signal degradation) • **ES-L**—Errored seconds (near-end line) • **SES-L**—Severely errored seconds (near-end line) • **UAS-L**—Unavailable seconds (near-end line) • **ES-LFE**—Errored seconds (far-end line) • **SES-LFE**—Severely errored seconds (far-end line) • **UAS-LFE**—Unavailable seconds (far-end line)	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 73: SONET/SDH show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
SONET path	Active alarms and defects, plus counts of specific SONET errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than**OK** indicates a problem. Subfields are: • **BIP-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) • **ES-PFE**—Errored seconds (far-end STS path) • **SES-PFE**—Severely errored seconds (far-end STS path) • **UAS-PFE**—Unavailable seconds (far-end STS path)	**extensive**
Received SONET overhead **Transmitted SONET overhead**	Values of the received and transmitted SONET overhead: • **C2**—Signal label. Allocated to identify the construction and content of the STS-level SPE and for PDI-P. • **F1**—Section user channel byte. This byte is set aside for the purposes of users. • **K1** and **K2**—These bytes are allocated for APS signaling for the protection of the multiplex section. • **J0**—Section trace. This byte is defined for STS-1 number 1 of an STS-*N* signal. Used to transmit a 1-byte fixed-length string or a 16-byte message so that a receiving terminal in a section can verify its continued connection to the intended transmitter. • **S1**—Synchronization status. The S1 byte is located in the first STS-1 of an STS-*N*. • **Z3** and **Z4**—Allocated for future use.	**extensive**
SDH alarms **SDH defects**	SDH media-specific defects that can prevent the interface from passing packets. When a defect persists for a certain period, 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 light the red or yellow alarm LED on the craft interface. See these fields for possible alarms and defects: **SDH PHY**, **SDH regenerator section**, **SDH multiplex section**, and **SDH path**.	All levels

 Copyright © 2018, Juniper Networks, Inc.

Table 73: SONET/SDH show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
SDH PHY	Active alarms and defects, plus counts of specific SDH errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **PLL Lock**—Phase-locked loop • **PHY Light**—Loss of optical signal	**extensive**
SDH regenerator section	Active alarms and defects, plus counts of specific SDH errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **RS-BIP8**—24-bit BIP for multiplex section overhead (B2 bytes) • **OOF**—Out of frame • **LOS**—Loss of signal • **LOF**—Loss of frame • **RS-ES**—Errored seconds (near-end regenerator section) • **RS-SES**—Severely errored seconds (near-end regenerator section) • **RS-SEFS**—Severely errored framing seconds (regenerator section)	**extensive**
SDH multiplex section	Active alarms and defects, plus counts of specific SDH errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **MS-BIP24**—8-bit BIP for high-order path overhead (B3 byte) • **MS-FEBE**—Far-end block error (multiplex section) • **MS-FERF**—Far-end remote fail (multiplex section) • **MS-AIS**—Alarm indication signal (multiplex section) • **BERR-SF**—Bit error rate fault (signal failure) • **BERR-SD**—Bit error rate defect (signal degradation) • **MS-ES**—Errored seconds (near-end multiplex section) • **MS-SES**—Severely errored seconds (near-end multiplex section) • **MS-UAS**—Unavailable seconds (near-end multiplex section) • **MS-ES-FE**—Errored seconds (far-end multiplex section) • **MS-SES-FE**—Severely errored seconds (far-end multiplex section) • **MS-UAS-FE**—Unavailable seconds (far-end multiplex section)	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 73: SONET/SDH show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
SDH path	Active alarms and defects, plus counts of specific SDH errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **HP-BIP8**—8-bit BIP for regenerator section overhead (B1 byte) • **HP-FEBE**—Far-end block error (high-order path) • **HP-LOP**—Loss of pointer (high-order path) • **HP-AIS**—High-order-path alarm indication signal • **HP-FERF**—Far-end remote fail (high-order path) • **HP-UNEQ**—Unequipped (high-order path) • **HP-PLM**—Payload label mismatch (high-order path) • **HP-ES**—Errored seconds (near-end high-order path) • **HP-SES**—Severely errored seconds (near-end high-order path) • **HP-UAS**—Unavailable seconds (near-end high-order path) • **HP-ES-FE**—Errored seconds (far-end high-order path) • **HP-SES-FE**—Severely errored seconds (far-end high-order path) • **HP-UAS-FE**—Unavailable seconds (far-end high-order path)	**extensive**
Received SDH overhead **Transmitted SDH overhead**	Values of the received and transmitted SONET overhead: • **C2**—Signal label. Allocated to identify the construction and content of the STS-level SPE and for PDI-P. • **F1**—Section user channel byte. This byte is set aside for the purposes of users. • **K1** and **K2**—These bytes are allocated for APS signaling for the protection of the multiplex section. • **J0**—Section trace. This byte is defined for STS-1 number 1 of an STS-*N* signal. Used to transmit a 1-byte fixed-length string or a 16-byte message so that a receiving terminal in a section can verify its continued connection to the intended transmitter. • **S1**—Synchronization status. The S1 byte is located in the first STS-1 of an STS-*N*. • **Z3** and **Z4**—Allocated for future use.	**extensive**
Received path trace **Transmitted path trace**	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**

Copyright © 2018, Juniper Networks, Inc.

Table 73: SONET/SDH show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
HDLC configuration	Information about the HDLC configuration. • **Policing bucket**—Configured state of the receiving policer. • **Shaping bucket**—Configured state of the transmitting shaper. • **Giant threshold**—Giant threshold programmed into the hardware. • **Runt threshold**—Runt threshold programmed into the hardware.	**extensive**
CoS information	Information about the CoS queue for the physical interface. • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number. • **PLP byte**—Packet Level Protocol byte.	**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 1110.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
PPP parameters	The PPP loopback clear timer value.	**extensive**

Table 73: SONET/SDH show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Shared interface	Provides the following information: • **shared with**—(RSD only) Indicates which PSD owns the logical shared interface. For example, **psd3**. • **peer interface**—(PSD only) Lists the logical tunnel interface that peers with the logical shared interface. For example, **ut-2/1/0.2**. • **tunnel token**—Specifies the receive (RX) and transmit (TX) tunnel tokens. For example, **Rx: 5.519, Tx: 13.514**.	All levels
Input packets	Number of packets received on the logical interface.	None specified
Output packets	Number of packets transmitted on the logical interface.	None specified
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	**detail extensive**
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	**detail extensive**
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	**detail extensive**
Protocol	Protocol family configured on the logical interface, such as **iso**, **inet6**, or **mpls**.	**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**
Multilink bundle	(If the logical interface is configured as part of a multilink bundle.) Interface name for the multilink bundle.	**detail extensive** none
AS bundle	(If the logical interface is configured as part of an aggregated SONET bundle.) AS bundle number.	**detail extensive**
MTU	MTU size on the logical interface.	**detail extensive none**
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route table	Routing table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive**
Flags	Information about the protocol family flags. Possible values are described in the “Family Flags” section under “Common Output Fields Description” on page 1110.	**detail extensive** none

 Copyright © 2018, Juniper Networks, Inc.

Table 73: SONET/SDH show interfaces 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 1110.	**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 interface.	**detail extensive** none
DLCI	(Frame Relay) DLCI number of the logical interface. The following DLCI information is displayed: **Flags**, **Total down time**, **Last down**, and **Traffic statistics**. **Flags** is one or more of the following: • **Active**—Set when the link is active and the DTE and DCE are exchanging information. • **Down**—Set when the link is active, but no information is received from the DCE. • **Unconfigured**—Set when the corresponding DLCI in the DCE is not configured. • **Configured**—Set when the corresponding DLCI in the DCE is configured. • **Dce-configured**—Displayed when the command is issued from the DTE.	**detail extensive**
DLCI statistics	(Frame Relay) Data-link connection identifier (DLCI) statistics. • **Active DLCI**—Number of active DLCIs. • **Inactive DLCI**—Number of inactive DLCIs.	**detail extensive** none

Sample Output

show interfaces (SDH Mode, PPP)

```
user@host> show interfaces so-0/0/0
Physical interface: so-0/0/0, Enabled, Physical link is Up
  Interface index: 149, SNMP ifIndex: 66
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SDH mode, Speed: OC3,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 30 (00:00:07 ago), Output: 29 (00:00:05 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Closed
  CoS queues     : 4 supported, 4 maximum usable queues
  Last flapped   : 2006-03-24 13:20:56 PST (00:05:09 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SDH   alarms   : None
  SDH   defects  : None
```

Copyright © 2018, Juniper Networks, Inc.

```
  Logical interface so-0/0/0.0 (Index 66) (SNMP ifIndex 43)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
    Protocol inet, MTU: 4470
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.0.12.0/30, Local: 10.0.12.1, Broadcast: 10.0.12.3
    Protocol iso, MTU: 4470
      Flags: Protocol-Down
    Protocol mpls, MTU: 4458, Maximum labels: 3
      Flags: Protocol-Down, Is-Primary
```

show interfaces brief (SDH Mode, PPP)

```
user@host> show interfaces so-0/0/0 brief
Physical interface: so-0/0/0, Enabled, Physical link is Up
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SDH mode, Speed: OC3,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 25 (00:00:01 ago), Output: 24 (00:00:04 ago)
  SDH   alarms   : None
  SDH   defects  : None

  Logical interface so-0/0/0.0
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
    inet  10.0.12.1/30
    iso
    mpls
```

show interfaces detail (SDH Mode, PPP)

```
user@host>  show interfaces so-0/0/0 detail
Physical interface: so-0/0/0, Enabled, Physical link is Up
  Interface index: 149, SNMP ifIndex: 66, Generation: 35
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SDH mode, Speed: OC3,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 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 : 33 (last seen 00:00:05 ago)
    Output: 32 (last sent 00:00:06 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Closed
  CoS queues     : 4 supported, 4 maximum usable queues
  Last flapped   : 2006-03-24 13:20:56 PST (00:05:38 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                  862                    0 bps
   Output bytes  :                 3592                   64 bps
   Input  packets:                   70                    0 pps
   Output packets:                  330                    0 pps
  Egress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
```

 Copyright © 2018, Juniper Networks, Inc.

```
    0 best-effort                   0                    0                    0

    1 expedited-fo                  0                    0                    0

    2 assured-forw                  0                    0                    0

    3 network-cont                329                  329                    0

  SDH   alarms   : None
  SDH   defects  : None

  Logical interface so-0/0/0.0 (Index 66) (SNMP ifIndex 43) (Generation 19)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
    Protocol inet, MTU: 4470, Generation: 48, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.0.12.0/30, Local: 10.0.12.1, Broadcast: 10.0.12.3,
        Generation: 48
    Protocol iso, MTU: 4470, Generation: 49, Route table: 0
      Flags: Protocol-Down
    Protocol mpls, MTU: 4458, Maximum labels: 3, Generation: 50, Route table: 0
      Flags: Protocol-Down, Is-Primary
```

show interfaces extensive (SDH Mode, PPP)

```
user@host> show interfaces so-0/0/0 extensive
Physical interface: so-0/0/0, Enabled, Physical link is Up
  Interface index: 149, SNMP ifIndex: 66, Generation: 35
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SDH mode, Speed: OC3,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 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 : 36 (last seen 00:00:01 ago)
    Output: 35 (last sent 00:00:10 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Closed
  CoS queues     : 4 supported, 4 maximum usable queues
  Last flapped   : 2006-03-24 13:20:56 PST (00:06:08 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                  922                    0 bps
   Output bytes  :                 3850                   64 bps
   Input  packets:                   75                    0 pps
   Output packets:                  356                    0 pps
  Label-switched interface (LSI) traffic statistics:
   Input  bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Bucket drops: 0, Policed discards: 218, L3 incompletes: 0,
    L2 channel errors: 0, L2 mismatch timeouts: 2, HS link CRC errors: 0,
    HS link FIFO overflows: 0
  Output errors:
    Carrier transitions: 3, Errors: 0, Drops: 0, Aged packets: 0,
```

```
  HS link FIFO underflows: 0, MTU errors: 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                 354                  354                    0

SDH   alarms   : None
SDH   defects  : None
SDH PHY:              Seconds        Count  State
  PLL Lock                  0            0  OK
  PHY Light                 2            1  OK
SDH regenerator section:
  RS-BIP8                   0            0
  OOF                       3            8  OK
  LOS                       3            2  OK
  LOF                       3            2  OK
  RS-ES                     3
  RS-SES                    3
  RS-SEFS                   3
SDH multiplex section:
  MS-BIP24                  0            0
  MS-FEBE                   0            0
  MS-FERF                   3            2  OK
  MS-AIS                    2            1  OK
  BERR-SF                   0            0  OK
  BERR-SD                   0            0  OK
  MS-ES                     3
  MS-SES                    3
  MS-UAS                    0
  MS-SES-FE                 3
  MS-UAS-FE                 0
SDH path:
  HP-BIP8                   0            0
  HP-FEBE                   0            0
  HP-LOP                    1            1  OK
  HP-AIS                    2            1  OK
  HP-FERF                   3            2  OK
  HP-UNEQ                   0            0  OK
  HP-PLM                    1            1  OK
  HP-ES                     3
  HP-SES                    3
  HP-UAS                    0
  HP-ES-FE                  3
  HP-SES-FE                 3
  HP-UAS-FE                 0
Received SDH overhead:
  F1      : 0x00, J0      : 0x00, K1      : 0x00, K2      : 0x00
  S1      : 0x00, C2      : 0xcf, C2(cmp) : 0xcf, F2      : 0x00
  Z3      : 0x00, Z4      : 0x00, S1(cmp) : 0x00
Transmitted SDH overhead:
  F1      : 0x00, J0      : 0x01, K1      : 0x00, K2      : 0x00
  S1      : 0x00, C2      : 0xcf, F2      : 0x00, Z3      : 0x00
  Z4      : 0x00
Received path trace: R2 so-0/0/0
  52 32 20 73 6f 2d 30 2f 30 2f 30 00 00 00 00 00   R2 so-0/0/0.....
```

 Copyright © 2018, Juniper Networks, Inc.

```
  Transmitted path trace: R1 so-0/0/0
    52 31 20 73 6f 2d 30 2f 30 2f 30 00 00 00 00 00   R1 so-0/0/0.....
  HDLC configuration:
    Policing bucket: Disabled
    Shaping bucket : Disabled
    Giant threshold: 4484, Runt threshold: 3
  Packet Forwarding Engine configuration:
    Destination slot: 0, PLP byte: 1 (0x00)
  CoS information:
    CoS transmit queue    Bandwidth      Buffer Priority   Limit
                     %         bps   %   usec
    0 best-effort     95  147744000  95      0      low    none
    3 network-control  5    7776000   5      0      low    none

  Logical interface so-0/0/0.0 (Index 66) (SNMP ifIndex 43) (Generation 19)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
    PPP parameters:
      PPP loopback clear timer: 3 sec
    Protocol inet, MTU: 4470, Generation: 48, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.0.12.0/30, Local: 10.0.12.1, Broadcast: 10.0.12.3,
        Generation: 48
    Protocol iso, MTU: 4470, Generation: 49, Route table: 0
      Flags: Protocol-Down
    Protocol mpls, MTU: 4458, Maximum labels: 3, Generation: 50, Route table: 0
      Flags: Protocol-Down, Is-Primary
  MS-ES-FE                  3
```

show interfaces brief (SONET Mode, Frame Relay)

```
user@host> show interfaces so-0/0/0 brief
Physical interface: so-0/0/0, Enabled, Physical link is Up
  Link-level type: Frame-Relay, MTU: 4474, Clocking: Internal, SONET mode,
  Speed: OC3, Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Link-Layer-Down Point-To-Point SNMP-Traps Internal: 0x4000
  Link flags     : Keepalives DTE
  ANSI LMI settings: n391dte 6, n392dte 3, n393dte 4, t391dte 10 seconds
  LMI: Input: 29 (00:00:02 ago), Output: 28 (00:00:01 ago)
  SONET alarms   : None
  SONET defects  : None

  Logical interface so-0/0/0.0
    Flags: Device-Down Point-To-Point SNMP-Traps Encapsulation: FR-NLPID
    inet  10.0.12.1        --> 10.0.12.2
    iso
    mpls
    DLCI 16
      Flags: Down, DCE-Unconfigured
      Total down time: 00:04:12 sec, Last down: 00:04:12 ago
```

show interfaces (SONET Mode, Frame Relay)

```
user@host> show interfaces so-0/0/0
Physical interface: so-0/0/0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 66
  Link-level type: Frame-Relay, MTU: 4474, Clocking: Internal, SONET mode,
  Speed: OC3, Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
```

Copyright © 2018, Juniper Networks, Inc.

```
  Interface flags: Link-Layer-Down Point-To-Point SNMP-Traps Internal: 0x4000
  Link flags     : Keepalives DTE
  ANSI LMI settings: n391dte 6, n392dte 3, n393dte 4, t391dte 10 seconds
  LMI: Input: 23 (00:00:05 ago), Output: 22 (00:00:03 ago)
  DTE statistics:
    Enquiries sent                       : 19
    Full enquiries sent                  : 3
    Enquiry responses received           : 20
    Full enquiry responses received      : 3
  DCE statistics:
    Enquiries received                   : 0
    Full enquiries received              : 0
    Enquiry responses sent               : 0
    Full enquiry responses sent          : 0
  Common statistics:
    Unknown messages received            : 0
    Asynchronous updates received        : 0
    Out-of-sequence packets received     : 0
    Keepalive responses timedout         : 1
  CoS queues     : 4 supported, 4 maximum usable queues
  Last flapped   : 2006-03-06 11:53:20 PST (3d 03:09 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 56 bps (0 pps)
  SONET alarms   : None
  SONET defects  : None

  Logical interface so-0/0/0.0 (Index 79) (SNMP ifIndex 43)
    Flags: Device-Down Point-To-Point SNMP-Traps Encapsulation: FR-NLPID
    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: 4470
      Flags: None
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.0.12.2, Local: 10.0.12.1
    Protocol iso, MTU: 4470
      Flags: None
    Protocol mpls, MTU: 4450, Maximum labels: 3
    DLCI 16
      Flags: Down, DCE-Unconfigured
      Total down time: 00:03:11 sec, Last down: 00:03:11 ago
        Input packets : 0
        Output packets: 0
  DLCI statistics:
    Active DLCI  :0  Inactive DLCI  :1
```

show interfaces detail (SONET Mode, Frame Relay)

```
user@host> show interfaces so-0/0/0 detail
Physical interface: so-0/0/0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 66, Generation: 11
  Link-level type: Frame-Relay, MTU: 4474, Clocking: Internal, SONET mode,
  Speed: OC3, Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Link-Layer-Down Point-To-Point SNMP-Traps Internal: 0x4000
  Link flags     : Keepalives DTE
  Hold-times     : Up 0 ms, Down 0 ms
  ANSI LMI settings: n391dte 6, n392dte 3, n393dte 4, t391dte 10 seconds
  LMI statistics:
    Input : 33 (last seen 00:00:09 ago)
    Output: 32 (last sent 00:00:01 ago)
  DTE statistics:
```

 Copyright © 2018, Juniper Networks, Inc.

```
  Enquiries sent                      : 27
  Full enquiries sent                 : 5
  Enquiry responses received          : 28
  Full enquiry responses received     : 5
DCE statistics:
  Enquiries received                  : 0
  Full enquiries received             : 0
  Enquiry responses sent              : 0
  Full enquiry responses sent         : 0
Common statistics:
  Unknown messages received           : 0
  Asynchronous updates received       : 0
  Out-of-sequence packets received    : 0
  Keepalive responses timedout        : 1
CoS queues     : 4 supported, 4 maximum usable queues
Last flapped   : 2006-03-06 11:53:20 PST (3d 03:10 ago)
Statistics last cleared: Never
Traffic statistics:
 Input  bytes  :             495368                    0 bps
 Output bytes  :            2765014                   56 bps
 Input  packets:              41165                    0 pps
 Output packets:             133530                    0 pps
Egress queues: 4 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets

  0 best-effort                  18                   18                    0

  1 expedited-fo                  0                    0                    0

  2 assured-forw                  0                    0                    0

  3 network-cont             133506               133506                    0

SONET alarms   : None
SONET defects  : None
Logical interface so-0/0/0.0 (Index 79) (SNMP ifIndex 43) (Generation 28)
  Flags: Device-Down Point-To-Point SNMP-Traps Encapsulation: FR-NLPID
  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: 4470, Generation: 49, Route table: 0
    Flags: None
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 10.0.12.2, Local: 10.0.12.1, Broadcast: Unspecified,
      Generation: 61
  Protocol iso, MTU: 4470, Generation: 50, Route table: 0
    Flags: None
  Protocol mpls, MTU: 4450, Maximum labels: 3, Generation: 51, Route table: 0
  DLCI 16
```

Copyright © 2018, Juniper Networks, Inc.

```
    Flags: Down, DCE-Unconfigured
    Total down time: 00:04:54 sec, Last down: 00:04:54 ago
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
DLCI statistics:
  Active DLCI  :0  Inactive DLCI  :1
```

show interfaces extensive (SONET Mode, Frame Relay)

```
user@host> show interfaces so-0/0/0 extensive
Physical interface: so-0/0/0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 66, Generation: 11
  Link-level type: Frame-Relay, MTU: 4474, Clocking: Internal, SONET mode,
  Speed: OC3, Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Link-Layer-Down Point-To-Point SNMP-Traps Internal: 0x4000
  Link flags     : Keepalives DTE
  Hold-times     : Up 0 ms, Down 0 ms
  ANSI LMI settings: n391dte 6, n392dte 3, n393dte 4, t391dte 10 seconds
  LMI statistics:
    Input : 39 (last seen 00:00:02 ago)
    Output: 36 (last sent 00:00:07 ago)
  DTE statistics:
    Enquiries sent                       : 30
    Full enquiries sent                  : 6
    Enquiry responses received           : 33
    Full enquiry responses received      : 6
  DCE statistics:
    Enquiries received                   : 0
    Full enquiries received              : 0
    Enquiry responses sent               : 0
    Full enquiry responses sent          : 0
  Common statistics:
    Unknown messages received            : 0
    Asynchronous updates received        : 0
    Out-of-sequence packets received     : 0
    Keepalive responses timedout         : 1
  CoS queues     : 4 supported, 4 maximum usable queues
  Last flapped   : 2006-03-06 11:53:20 PST (3d 03:11 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              495452                   56 bps
   Output bytes  :             2765074                    0 bps
   Input  packets:               41171                    0 pps
   Output packets:              133534                    0 pps
  Label-switched interface (LSI) traffic statistics:
   Input  bytes  :                   0                    0 bps
   Input  packets:                   0                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Bucket drops: 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
  Output errors:
    Carrier transitions: 3, Errors: 0, Drops: 0, Aged packets: 0,
    HS link FIFO underflows: 0, MTU errors: 0
  Egress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
```

 Copyright © 2018, Juniper Networks, Inc.

```
    0 best-effort                  18                    18                     0

    1 expedited-fo                  0                     0                     0

    2 assured-forw                  0                     0                     0

    3 network-cont             133510                133510                     0

  SONET alarms   : None
  SONET defects  : None
  SONET PHY:            Seconds        Count  State
    PLL Lock                  0            0  OK
    PHY Light                60            1  OK
  SONET section:
    BIP-B1                    0            0
    SEF                     108          158  OK
    LOS                     108            2  OK
    LOF                     108            2  OK
    ES-S                    108
    SES-S                   108
    SEFS-S                  108
  SONET line:
    BIP-B2                    0            0
    REI-L                     0            0
    RDI-L                     1            1  OK
    AIS-L                   107            1  OK
    BERR-SF                   0            0  OK
    BERR-SD                  44            2  OK
    ES-L                    108
    SES-L                   108
    UAS-L                    97
    ES-LFE                    1
    SES-LFE                   1
    UAS-LFE                   0
SONET path:
    BIP-B3                    0            0
    REI-P                     0            0
    LOP-P                     1            1  OK
    AIS-P                   107            1  OK
    RDI-P                     1            1  OK
    UNEQ-P                    0            0  OK
    PLM-P                     1            1  OK
    ES-P                    108
    SES-P                   108
    UAS-P                    97
    ES-PFE                    1
    SES-PFE                   1
    UAS-PFE                   0
  Received SONET overhead:
    F1      : 0x00, J0      : 0x00, K1      : 0x00, K2      : 0x00
    S1      : 0x00, C2      : 0xcf, C2(cmp) : 0xcf, F2      : 0x00
    Z3      : 0x00, Z4      : 0x00, S1(cmp) : 0x00
  Transmitted SONET overhead:
    F1      : 0x00, J0      : 0x01, K1      : 0x00, K2      : 0x00
    S1      : 0x00, C2      : 0xcf, F2      : 0x00, Z3      : 0x00
    Z4      : 0x00
  Received path trace: R2 so-0/0/0
    52 32 20 73 6f 2d 30 2f 30 2f 30 00 00 00 00 00   R2 so-0/0/0.....
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
```

Copyright © 2018, Juniper Networks, Inc.

```
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 0d 0a ................
  Transmitted path trace: R1 so-0/0/0
    52 31 20 73 6f 2d 30 2f 30 2f 30 00 00 00 00 00   R1 so-0/0/0.....
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
  HDLC configuration:
    Policing bucket: Disabled
    Shaping bucket : Disabled
    Giant threshold: 4484, Runt threshold: 3
  Packet Forwarding Engine configuration:
    Destination slot: 0, PLP byte: 1 (0x00)
  CoS information:
    CoS transmit queue        Bandwidth              Buffer  Priority   Limit
                         %          bps     %           usec
    0 best-effort       95    147744000    95              0       low    none
    3 network-control    5      7776000     5              0       low    none

  Logical interface so-0/0/0.0 (Index 79) (SNMP ifIndex 43) (Generation 28)
    Flags: Device-Down Point-To-Point SNMP-Traps Encapsulation: FR-NLPID
    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: 4470, Generation: 49, Route table: 0
      Flags: None
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.0.12.2, Local: 10.0.12.1, Broadcast: Unspecified,
        Generation: 61
    Protocol iso, MTU: 4470, Generation: 50, Route table: 0
      Flags: None
    Protocol mpls, MTU: 4450, Maximum labels: 3, Generation: 51, Route table: 0
    DLCI 16
      Flags: Down, DCE-Unconfigured
      Total down time: 00:05:42 sec, Last down: 00:05:42 ago
      Traffic statistics:
       Input  bytes  :                    0
       Output bytes  :                    0
       Input  packets:                    0
       Output packets:                    0
  DLCI statistics:
    Active DLCI   :0  Inactive DLCI   :1
```

show interfaces extensive (OC768-over-4xOC192 Mode)

```
user@host> show interfaces so-7/0/0 extensive
Physical interface: so-7/0/0, Enabled, Physical link is Up
  Interface index: 163, SNMP ifIndex: 23, Generation: 186
  Link-level type: Cisco-HDLC, MTU: 4474, Clocking: Internal, SONET mode, Speed:
 OC768,
```

 Copyright © 2018, Juniper Networks, Inc.

```
  Loopback: Local, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Link flags     : No-Keepalives
  Hold-times     : Up 0 ms, Down 0 ms
  CoS queues     : 8 supported, 8 maximum usable queues
  Last flapped   : 2006-01-13 10:43:39 PST (01:05:33 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :               76992                  200 bps
   Output bytes  :               83707                  216 bps
   Input  packets:                1343                    0 pps
   Output packets:                1343                    0 pps
  Input errors:
    Errors: 0, Drops: 3885, Framing errors: 68154624, Runts: 0, Giants: 0, Bucket
 drops: 0,
    Policed discards: 0, L3 incompletes: 95040248, L2 channel errors: 0, L2
mismatch timeouts: 0,
    HS link CRC errors: 0, HS link FIFO overflows: 30742070
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, Aged packets: 0, HS link FIFO
underflows: 0,
    MTU errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                    2                    2                    0

    1 expedited-fo                   0                    0                    0

    2 assured-forw                   0                    0                    0

    3 network-cont                1341                 1341                    0

  SONET alarms   : None
  SONET defects  : None
  Link : 0
  SONET alarms   : None
  SONET defects  : None
  SONET PHY:          Seconds        Count  State
    PLL Lock                0            0  OK
    PHY Light               0            0  OK
  SONET section:
    BIP-B1                  0            0
    SEF                     2            1  OK
    LOS                     0            0  OK
    LOF                     3            2  OK
    ES-S                    2
    SES-S                   2
    SEFS-S                  2
  SONET line:
    BIP-B2                  0            0
    REI-L                   0            0
    RDI-L                   1            1  OK
    AIS-L                   2            1  OK
    BERR-SF                 0            0  OK
    BERR-SD                 0            0  OK
    ES-L                    3
    SES-L                   3
    UAS-L                   0
    ES-LFE                  1
```

```
  SES-LFE                  1
  UAS-LFE                  0
SONET path:
  BIP-B3                   0            0
  REI-P                    0            0
  LOP-P                    0            0  OK
  AIS-P                    2            1  OK
  RDI-P                    0            0  OK
  UNEQ-P                   0            0  OK
  PLM-P                    0            0  OK
  ES-P                     3
  SES-P                    3
  UAS-P                    0
  ES-PFE                   0
  SES-PFE                  0
  UAS-PFE                  0
Payload pointer:
  Current pointer             :  522
  Pointer increment count     :  0
  Pointer decrement count     :  0
  New pointer NDF count       :  0
Received SONET overhead:
  F1      : 0x00, J0      : 0x01, K1      : 0x00, K2      : 0x00
  S1      : 0x00, C2      : 0xcf, C2(cmp) : 0xcf, F2      : 0x00
  Z3      : 0x00, Z4      : 0x00, S1(cmp) : 0x00
Transmitted SONET overhead:
  F1      : 0x00, J0      : 0x01, K1      : 0x00, K2      : 0x00
  S1      : 0x00, C2      : 0xcf, F2      : 0x00, Z3      : 0x00
  Z4      : 0x00
Received path trace: fold so-7/0/0
  66 6f 6c 64 20 73 6f 2d 37 2f 30 2f 30 00 00 00   fold so-7/0/0...
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 0d 0a ................
Transmitted path trace: fold so-7/0/0
  66 6f 6c 64 20 73 6f 2d 37 2f 30 2f 30 00 00 00   fold so-7/0/0...
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
Link : 1
SONET alarms   : None
SONET defects  : None
SONET PHY:            Seconds        Count  State
  PLL Lock                 0            0  OK
  PHY Light                0            0  OK
SONET section:
  BIP-B1                   0            0
  SEF                      2            1  OK
  LOS                      0            0  OK
  LOF                      3            2  OK
  ES-S                     2
  SES-S                    2
  SEFS-S                   2
SONET line:
  BIP-B2                   0            0
  REI-L                    0            0
  RDI-L                    0            0  OK
  AIS-L                    2            1  OK
  BERR-SF                  0            0  OK
  BERR-SD                  0            0  OK
  ES-L                     3
```

 Copyright © 2018, Juniper Networks, Inc.

```
      SES-L                     3
      UAS-L                     0
      ES-LFE                    0
      SES-LFE                   0
      UAS-LFE                   0
    SONET path:
      BIP-B3                    0             0
      REI-P                     0             0
      LOP-P                     0             0  OK
      AIS-P                     2             1  OK
      RDI-P                     0             0  OK
      UNEQ-P                    0             0  OK
      PLM-P                     0             0  OK
      ES-P                      3
      SES-P                     3
      UAS-P                     0
      ES-PFE                    0
      SES-PFE                   0
      UAS-PFE                   0
    Payload pointer:
      Current pointer               :  522
      Pointer increment count       :  0
      Pointer decrement count       :  0
      New pointer NDF count         :  0
    Received SONET overhead:
      F1       : 0x00, J0       : 0x01, K1       : 0x00, K2       : 0x00
      S1       : 0x00, C2       : 0xcf, C2(cmp)  : 0xcf, F2       : 0x00
      Z3       : 0x00, Z4       : 0x00, S1(cmp)  : 0x00
    Transmitted SONET overhead:
      F1       : 0x00, J0       : 0x01, K1       : 0x00, K2       : 0x00
      S1       : 0x00, C2       : 0xcf, F2       : 0x00, Z3       : 0x00
      Z4       : 0x00
    Received path trace: fold so-7/0/0
      66 6f 6c 64 20 73 6f 2d 37 2f 30 2f 30 00 00 00   fold so-7/0/0...
      00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
      00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
      00 00 00 00 00 00 00 00 00 00 00 00 00 00 0d 0a ................
    Transmitted path trace: fold so-7/0/0
      66 6f 6c 64 20 73 6f 2d 37 2f 30 2f 30 00 00 00   fold so-7/0/0...
      00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
      00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
      00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
...
```

show interfaces detail (IPv6 Tracking)

```
user@host> show interfaces so-0/2/0 detail
Physical interface: so-0/2/0, Enabled, Physical link is Up
   Interface index: 130, SNMP ifIndex: 26, Generation: 131
   Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode,
Speed: OC3,
   Loopback: None, FCS: 16, Payload scrambler: Enabled
   Device flags   : Present Running
   Interface flags: Point-To-Point SNMP-Traps 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 : 7 (last seen 00:00:01 ago)
     Output: 6 (last sent 00:00:08 ago)
   LCP state: Opened
```

```
NCP state: inet: Not-configured, inet6: Opened, iso: Not- configured, mpls:
Not-configured
CHAP state: Closed
PAP state: Closed
CoS queues     : 4 supported, 4 maximum usable queues
Last flapped   : 2007-11-29 08:45:47 PST (1d 03:44 ago)
Statistics last cleared: Never
Traffic statistics:
 Input  bytes  :              7407782                   40 bps
 Output bytes  :              7307322                   48 bps
 Input  packets:               107570                    0 pps
 Output packets:               108893                    0 pps
 IPv6 transit statistics:
  Input  bytes  :               57328
  Output bytes  :               57400
  Input  packets:                1024
  Output packets:                1025
Egress queues: 4 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets   Dropped packets
  0 best-effort       1191             1191                 0
  1 expedited-fo      0                0                    0
  2 assured-forw      0                0                    0
  3 network-cont      107700           107700               0
SONET alarms   : None
SONET defects  : None

Logical interface so-0/2/0.0 (Index 70) (SNMP ifIndex 47) (Generation 231)
  Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
  Protocol inet6, MTU: 4470, Generation: 433, Route table: 0
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
      Destination:  2001:db8::2:1/32, Local:  2001:db8::2:2,
      Broadcast: Unspecified, Generation: 683
    Addresses, Flags: Is-Preferred
      Destination:  2001:db8::1:2, Local: 2001:db8::1:3,
      Broadcast: Unspecified, Generation: 684
```

show interfaces (Shared Interface)

```
user@rsd1> show interfaces so-7/2/0
Physical interface: so-7/2/0, Enabled, Physical link is Down
  Interface index: 128, SNMP ifIndex: 109
  Link-level type: Frame-Relay, MTU: 4474, Clocking: Internal, SONET mode,
  Speed: OC192, Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down Point-To-Point SNMP-Traps Internal: 0x4000
  Shared-interface  : Owner
  Link flags     : No-Keepalives DTE
  ANSI LMI settings: n391dte 6, n392dte 3, n393dte 4, t391dte 10 seconds
  LMI: Input: 0 (never), Output: 0 (never)
  DTE statistics:
    Enquiries sent                       : 0
    Full enquiries sent                  : 0
    Enquiry responses received           : 0
    Full enquiry responses received      : 0
  DCE statistics:
    Enquiries received                   : 0
    Full enquiries received              : 0
    Enquiry responses sent               : 0
    Full enquiry responses sent          : 0
  Common statistics:
```

 Copyright © 2018, Juniper Networks, Inc.

```
  Unknown messages received           : 0
  Asynchronous updates received       : 0
  Out-of-sequence packets received    : 0
  Keepalive responses timedout        : 0
CoS queues     : 8 supported, 8 maximum usable queues
Last flapped   : 2008-08-11 10:51:51 PDT (1w1d 04:47 ago)
Input rate     : 0 bps (0 pps)
Output rate    : 0 bps (0 pps)
SONET alarms   : LOL, PLL
SONET defects  : LOL, PLL, LOF, SEF, AIS-L, AIS-P

Logical interface so-7/2/0.0 (Index 67) (SNMP ifIndex 117)
  Flags: Device-Down Point-To-Point SNMP-Traps 0x4000 Encapsulation: FR-NLPID
  Shared interface:
    Shared with: psd5
    Tunnel token: Rx: 2.517, Tx: 1.517
  Input packets : 0
  Output packets: 0
  DLCI 700
    Flags: Active
    Total down time: 00:01:09 sec, Last down: 284:58:21 ago
      Input packets : 0
      Output packets: 0
DLCI statistics:
  Active DLCI  :1  Inactive DLCI  :0
```

show interfaces (Serial)

Syntax

show interfaces *interface-type*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display status information about serial interfaces, including RS-232, RS-422/449, EIA-530, X.21, and V.35.

Options

interface-type—On M Series and T Series routers, the interface type is **se-*fpc/pic/port***.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

List of Sample Output

show interfaces (Serial, EIA-530) on page 1496
show interfaces brief (Serial, EIA-530) on page 1496
show interfaces detail (Serial, EIA-530) on page 1497
show interfaces extensive (Serial, EIA-530) on page 1497
show interfaces (Serial, V.35) on page 1498
show interfaces brief (Serial, V.35) on page 1499
show interfaces detail (Serial, V.35) on page 1499
show interfaces extensive (Serial, V.35) on page 1500
show interfaces statistics detail (RS 449) on page 1501

Output Fields

Table 74 on page 1490 lists the output fields for the **show interfaces** (Serial) command. Output fields are listed in the approximate order in which they appear.

Table 74: show interfaces (Serial) Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels

Copyright © 2018, Juniper Networks, Inc.

Table 74: 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 1110.	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 1110.	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 1110.	All levels
Link flags	Information about the link. Possible values are described in the “Link Flags” section under “Common Output Fields Description” on page 1110.	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. • **Interval** ***seconds***—Time between successive keepalive requests. The range of values, in seconds, is 10 to 32,767. The default value is 10. • **Up-count** ***number***—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** ***number***—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. • **Input:** ***number*** **(*hh*:*mm*:*ss* ago)**—Number of keepalive packets received by PPP and the time since the last keepalive packet was received. • **Output:** ***number*** **(*hh*:*mm*:*ss* ago)**—Number of keepalive packets sent by PPP and the time since the last keepalive packet was sent.	**brief** none

Copyright © 2018, Juniper Networks, Inc.

Table 74: show interfaces (Serial) Output Fields (continued)

Field Name	Field Description	Level of Output
Keepalive statistics	(PPP and HDLC) Information about keepalive packets. • **Input:** ***number*** **(last seen** ***hh:mm:ss*****ago)**—Number of keepalive packets received by PPP and the time since the last keepalive packet was received. • **Output:** ***number*****(last seen** ***hh:mm:ss*** **ago)**—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. • **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. • **Conf-ack-received**—Acknowledgement was received. • **Conf-ack-sent**—Acknowledgement was sent. • **Conf-req-sent**—Request was sent. • **Down**—NCP negotiation is incomplete (not yet completed or has failed). • **Not-configured**—NCP is not configured on the interface. • **Opened**—NCP negotiation is successful.	**detail extensive** none
CHAP state	(PPP) Displays the state of the Challenge Handshake Authentication Protocol (CHAP) during its transaction. • **Chap-Chal-received**—Challenge was received but response not yet sent. • **Chap-Chal-sent**—Challenge was sent. • **Chap-Resp-received**—Response was received for the challenge sent, but CHAP has not yet moved into the Success state. (Most likely with RADIUS authentication.) • **Chap-Resp-sent**—Response was sent for the challenge received. • **Closed**—CHAP authentication is incomplete. • **Failure**—CHAP authentication failed. • **Not-configured**—CHAP is not configured on the interface. • **Success**—CHAP authentication was successful.	**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

 Copyright © 2018, Juniper Networks, Inc.

Table 74: 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	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—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	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the 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**

Copyright © 2018, Juniper Networks, Inc.

Table 74: 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. • **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:
	• **Line protocol**—**eia530**, **eia530a**, **rs232**, **rs449**, **v.35**, or **x.21**..
	• **Resync history**—Information about resynchronization events:
	  ◦ **Sync loss count**—Number of times the synchronization was lost.
	• **Data signal**—(X.21 and V.35) Information about the data signal:
	  ◦ **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:
	  ◦ **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:
	  ◦ **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: • **Destination slot**—FPC slot number. • **PLP byte**—Packet Level Protocol byte.	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 74: show interfaces (Serial) Output Fields (continued)

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface: • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	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 1110.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
protocol-family	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 1110.	**detail extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 74: 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 1110.	**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
```

 Copyright © 2018, Juniper Networks, Inc.

```
    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
```

 Copyright © 2018, Juniper Networks, Inc.

```
  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
```

Copyright © 2018, Juniper Networks, Inc.

```
  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:
```


Copyright © 2018, Juniper Networks, Inc.

```
    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
```

 Copyright © 2018, Juniper Networks, Inc.

show interfaces (T1, E1, or DS)

Syntax

show interfaces *interface-type*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information Command introduced before Junos OS Release 7.4.

Description Display status information about the specified T1, E1, or DS interface.

Options

interface-type—On ACX Series, M Series, MX Series, and T Series routers, the T1 interface type is **t1-*fpc/pic/port***, whereas the E1 interface type is **e1-*fpc/pic/port***, and DS interface type is **ds-*fpc/pic/port*:<channel>**.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level view

Related Documentation

• Understanding Interfaces on ACX Series Universal Metro Routers on page 11

List of Sample Output

Output Fields Table 75 on page 1504 lists the output fields for the **show interfaces** (T1 or E1) command. Output fields are listed in the approximate order in which they appear.

Table 75: T1 or E1 show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under “Common Output Fields Description” on page 1110.	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**
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: **Internal** or **External**.	All levels
Speed	Speed at which the interface is running.	All levels
Loopback	Whether loopback is enabled and the type of loopback (**local** or **remote**).	All levels
FCS	Frame check sequence on the interface (either **16** or **32**). The default is **16** bits.	All levels
Framing	Physical layer framing format used for the E1 interface on the link: **G704**, **G704-NO-CRC4**, or **Unframed**. The default is **G704**. Physical layer framing format used for the T1 interface on the link: **SF** and **ESF**. The default is **ESF**.	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 1110.	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 1110.	All levels
Link flags	Information about the link. Possible values are described in the “Link Flags” section under “Common Output Fields Description” on page 1110.	All levels
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 75: T1 or E1 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
IMA Link alarms	Current active IMA link alarms, including the following: • **LIF** • **LODS** • **RFI-IMA** • **Tx-Mis-Connected** • **Tx-Unusable-FE** • **Rx-Unusable-FE** • **Link Fault**	**detail extensive** none
IMA Link defects	Current active IMA link defects, including the following: • **LIF** • **LODS** • **RFI-IMA** • **Tx-Mis-Connected** • **Tx-Unusable-FE** • **Rx-Unusable-FE** • **Link Fault**	**detail extensive** none

IMA Link state	Current active IMA link status, including the following:
	• **Line:** synchronized or not synchronized
	• **Near end:**—Status of near-end receive and transmit links
	  ◦ **Rx:** Usable or Unusable
	  ◦ **Tx:** Usable or Unusable
	• **Far end:**—Status of far-end receive and transmit links
	  ◦ **Rx:** Usable or Unusable
	  ◦ **Tx:** Usable or Unusable
		detail extensive none

Copyright © 2018, Juniper Networks, Inc.

Table 75: T1 or E1 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
IMA link media	IMA Link Media Status, which provides the seconds and count state for the following link media parameters: • LIF • LODS • Err-ICP • IV • Rx-FC • Tx-FC • FE-Defects • FE-Rx-FC • FE-Tx-FC • Rx-ICP • Rx-Stuff • Tx-ICP • Tx-Stuff • Rx-SES • Rx-UAS • Rx-UUS • Tx-UUS • FE-Rx-SES • FE-Rx-UAS • FE-Rx-UUS • FE-Tx-UUS	**detail extensive** none
Keepalive settings	(PPP and HDLC) Configured settings for keepalives. • **interval** ***seconds***—The time in seconds between successive keepalive requests. The range is **10** seconds through **32,767** seconds, with a default of **10** seconds. • **down-count** ***number***—The number of keepalive packets a destination must fail to receive before the network takes a link down. The range is **1** through **255**, with a default of **3**. • **up-count** ***number***—The number of keepalive packets a destination must receive to change a link's status from down to up. The range is **1** through **255**, with a default of **1**.	**detail extensive** none

Keepalive statistics	(PPP and HDLC) Information about keepalive packets. (When no level of output is specified, the word **statistics** is not part of the field name and the **last seen** text is not displayed.)
	• **Input**—Number of keepalive packets received by PPP.
	  ◦ **(last seen 00:00:00 ago)**—Time since the last keepalive packet was received, in the format ***hh:mm:ss***.
	• **Output**—Number of keepalive packets sent by PPP and how long ago the last keepalive packets were sent and received.
	  ◦ **(last seen 00:00:00 ago)**—Time since the last keepalive packet was sent, in the format ***hh:mm:ss***.
		detail extensive none

Copyright © 2018, Juniper Networks, Inc.

Table 75: T1 or E1 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
LMI settings	(Frame Relay) Settings for Local Management Interface (LMI) which can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is **(ANSI or ITU) LMI settings: *value*, *value*... *xx*** seconds, where *value* can be: • **n391dte**—DTE full status polling interval (1–255) • **n392dce**—DCE error threshold (1–10) • **n392dte**—DTE error threshold (1–10) • **n393dce**—DCE monitored event count (1–10) • **n393dte**—DTE monitored event count (1–10) • **t391dte**—DTE polling timer (5–30 seconds) • **t392dce**—DCE polling verification timer (5–30 seconds)	**detail extensive** none
LMI	(Frame Relay) Local Management Interface (LMI) packet statistics: • **Input**—Number of packets coming in on the interface (*nn*) and how much time has passed since the last packet arrived. The format is **Input: *nn* (last seen *hh*:*mm*:*ss* ago)**. • **Output**—Number of packets sent out on the interface (*nn*) and how much time has passed since the last packet was sent. The format is **Output: *nn* (last sent *hh*:*mm*:*ss* ago)**.	**detail extensive** none
DTE statistics	(Frame Relay) Statistics about messages transmitted from the data terminal equipment (DTE) to the data communications equipment (DCE): • **Enquiries sent**—Number of link status enquiries sent from the DTE to the DCE. • **Full enquiries sent**—Number of full enquiries sent from the DTE to the DCE. • **Enquiry responses received**—Number of enquiry responses received by the DTE from the DCE. • **Full enquiry responses received**—Number of full enquiry responses sent from the DTE to the DCE.	**detail extensive** none
DCE statistics	(Frame Relay) Statistics about messages transmitted from the DCE to the DTE: • **Enquiries received**—Number of enquiries received by the DCE from the DTE. • **Full enquiries received**—Number of full enquiries received by the DCE from the DTE. • **Enquiry responses sent**—Number of enquiry responses sent from the DCE to the DTE. • **Full enquiry responses sent**—Number of full enquiry responses sent from the DCE to the DTE.	**detail extensive** none

Copyright © 2018, Juniper Networks, Inc.

Table 75: T1 or E1 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Common statistics	(Frame Relay) Statistics about messages sent between the DTE and the DCE: • **Unknown messages received**—Number of received packets that do not fall into any category. • **Asynchronous updates received**—Number of link status peer changes received. • **Out-of-sequence packets received**—Number of packets for which the sequence of the packets received is different from the expected sequence. • **Keepalive responses timedout**—Number of keepalive responses that timed out when no Local Management Interface (LMI) packet was reported for **n392dte** or **n393dce** intervals. (See **LMI settings**.)	**detail extensive** none
Nonmatching DCE-end DLCIs	(Frame Relay. Displayed only from the DTE.) Number of DLCIs configured from the DCE.	**detail extensive** none
LCP state	(PPP) Link Control Protocol state. • **Conf-ack-received**—Acknowledgement was received. • **Conf-ack-sent**—Acknowledgement was sent. • **Conf-req-sent**—Request was sent. • **Down**—LCP negotiation is incomplete (not yet completed or has failed). • **Not configured**—LCP is not configured on the interface. • **Opened**—LCP negotiation is successful.	**detail extensive** none
NCP state	(PPP) Network Control Protocol state. • **Conf-ack-received**—Acknowledgement was received. • **Conf-ack-sent**—Acknowledgement was sent. • **Conf-req-sent**—Request was sent. • **Down**—NCP negotiation is incomplete (not yet completed or has failed). • **Not configured**—NCP is not configured on the interface. • **Opened**—NCP negotiation is successful.	**detail extensive** none
CHAP state	(PPP) State of the Challenge Handshake Authentication Protocol (CHAP) during its transaction. • **Chap-Chal-received**—Challenge was received but response is 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. • **Down**—CHAP authentication is incomplete (not yet completed or has failed). • **Not-configured**—CHAP is not configured on the interface. • **Opened**—CHAP authentication was successful.	**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

Copyright © 2018, Juniper Networks, Inc.

Table 75: T1 or E1 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
CoS Queues	Number of CoS queues configured.	**detail extensive** none
Input rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **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. • **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. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • **SRAM errors**—Number of hardware errors that occurred in the static RAM (SRAM) on the PIC or PIM. If the value of this field increments, the PIC or PIM is malfunctioning. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 75: T1 or E1 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and 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. • **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. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**
Queue counters	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**
DS1 alarms **DS1 defects**	E1 media-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. The following lists all possible alarms and defects. For complete explanations of most of these alarms and defects, see *Bellcore Telcordia GR-499-CORE*. • **AIS**—Alarm indication signal. • **LOF**—Loss of frame. • **LOS**—Loss of signal. • **YLW**—Yellow alarm. Indicates errors at the remote site receiver.	**detail extensive** none

Copyright © 2018, Juniper Networks, Inc.

Table 75: T1 or E1 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
T1 media or **E1 media**	Counts of T1 or E1 media-specific errors. • **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. The T1 or E1 media-specific error types are: • **SEF**—Severely errored framing • **BEE**—Bit error • **AIS**—Alarm indication signal • **LOF**—Loss of frame • **LOS**—Loss of signal • **YELLOW**—Errors at the remote site receiver • **CRC Major**—Cyclic redundancy check major alarm threshold exceeded • **CRC Minor**—Cyclic redundancy check minor alarm threshold exceeded • **BPV**—Bipolar violation • **EXZ**—Excessive zeros • **LCV**—Line code violation • **PCV**—Pulse code violation • **CS**—Carrier state • **CRC**—Cyclic redundancy check • **FEBE**—Far-end block error (E1 only) • **LES**—Line error seconds • **ES**—Errored seconds • **BES**—Bursty errored seconds • **SES**—Severely errored seconds • **SEFS**—Severely errored framing seconds • **UAS**—Unavailable seconds	**extensive**
SAToP Configuration	Information about the SAToP configuration. • **payload-size**—Configure the payload size, in bytes (from 32 through 1024 bytes). • **idle-pattern**—An 8-bit hexadecimal pattern to replace TDM data in a lost packet (from 0 through 255). • **jitter-buffer-packets**—Number of packets in the jitter buffer (from 1 through 64 packets). • **jitter-buffer-latency**—Time delay in the jitter buffer (from 1 through 1000 milliseconds). • **excessive-packet-loss-rate**—Set packet loss options. The options are **groups**, **sample-period**, and **threshold**. • **sample-period**—Time required to calculate excessive packet loss rate (from 1000 through 65,535 milliseconds). • **threshold**—Percentile designating the threshold of excessive packet loss rate (1–100 percent).	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 75: T1 or E1 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
CESoPSN Configuration	Information about the CESoPSN configuration. • **packetization-latency**—Time required to create packets (from 1000 through 8000 microseconds). • **idle-pattern**—An 8-bit hexadecimal pattern to replace TDM data in a lost packet (from 0 through 255). • **jitter-buffer-packets**—Number of packets in the jitter buffer (from 1 through 64 packets). • **jitter-buffer-latency**—Time delay in the jitter buffer (from 1 through 1000 milliseconds). • **excessive-packet-loss-rate**—Set packet loss options. The options are **sample-period** and **threshold**. • **sample-period**—Time required to calculate excessive packet loss rate (from 1000 through 65,535 milliseconds). • **threshold**—Percentile designating the threshold of excessive packet loss rate (1–100 percent).	**extensive**
HDLC configuration	Information about the HDLC configuration. • **Policing bucket**—Configured state of the receiving policer. • **Shaping bucket**—Configured state of the transmitting shaper. • **Giant threshold**—Giant threshold programmed into the hardware. • **Runt threshold**—Runt threshold programmed into the hardware. • **Timeslots**—Time slots configured on the interface. • **Buildout**—(T1 only) Buildout setting: 0-132, 133-265, 266-398, 399-531, or 532-655 feet. • **Timeslots**—Configured time slots for the interface. • **Byte encoding**—(T1 only) Byte encoding used: **Nx64K** or **Nx56K**. • **Line encoding**—Line encoding used. For T1, the value can be **B8ZS** or **AMI**. For E1, the value is **HDB3**. • **Data inversion**—HDLC data inversion setting: **Enabled** or **Disabled**. • **Idle cycle flag**—Idle cycle flags. • **Start end flag**—Start and end flag.	**extensive**
DS1 BERT configuration	BERT (bit error rate test) checks the quality of the line. This output appears only when a BERT is run on the interface. • **BERT time period**—Configured total time period that the BERT is to run. • **Elapsed**—Actual time elapsed since the start of the BERT (in seconds). • **Induced error rate**—Configured rate at which the bit errors are induced in the BERT pattern. • **Algorithm**—Type of algorithm selected for the BERT.	**detail extensive** none
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number. • **PLP byte**—Packet Level Protocol byte.	**extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 75: T1 or E1 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface. • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	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 interface. Possible values are described in the "Interface Flags" section under "Common Output Fields Description" on page 1110.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Input packets	Number of packets received on the logical interface.	None specified
Output packets	Number of packets transmitted on the logical interface.	None specified
Traffic statistics	(Frame Relay) Number and rate of bytes and packets received and transmitted on the logical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Local statistics	(Frame Relay) Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.	**detail extensive**

Table 75: T1 or E1 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Transit statistics	(Frame Relay) Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. This counter normally stabilizes in less than 1 second.	**detail extensive**
Protocol	Protocol family configured on the logical interface, such as **iso**, **inet6**, **mlfr**, or **mpls**.	**detail extensive** none
Multilink bundle	Interface name for the multilink bundle, if configured.	**detail extensive** none
MTU	MTU size on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route table	Routing table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive**
Flags	Information about the protocol family flags. Possible values are described in the "Family Flags" section under "Common Output Fields Description" on page 1110.	**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 1110.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive** none
DLCI	(Frame Relay) DLCI number of the logical interface. The following DLCI information is displayed: **Flags**, **Total down time**, **Last down**, and **Traffic statistics** or (**Input packets, Output packets**). **Flags** can be one or more of the following: • **Active**—Set when the link is active and the DTE and DCE are exchanging information. • **Down**—Set when the link is active, but no information is received from the DCE. • **DCE-Unconfigured**—Set when the corresponding DLCI in the DCE is not configured. • **Configured**—Set when the corresponding DLCI in the DCE is configured. • **DCE-configured**—Displayed when the command is issued from the DTE.	**detail extensive** none
DLCI statistics	(Frame Relay) Data-link connection identifier (DLCI) statistics. • **Active DLCI**—Number of active DLCIs. • **Inactive DLCI**—Number of inactive DLCIs.	**detail extensive** none

 Copyright © 2018, Juniper Networks, Inc.

Table 75: T1 or E1 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
CE Info	Information related to the circuit emulation statistics. • **CE Tx**—Number of transmitted packets and bytes (TDM to PSN flow). • **CE Rx**—Number of received packets and bytes and forward bytes (PSN to TDM flow). • **CE Rx Forwarded**—Number of forwarded bytes. • **CE Strayed**—Number of stray packets. • **CE Lost**—Number of lost packets. • **CE Malformed**—Number of malformed packets • **CE Misinserted**—Number of misinserted packets. • **CE AIS dropped**—Number of dropped bytes due to buffer overrun (PSN to TDM). • **CE Dropped**—Number of dropped packets during resynchronization • **CE Overrun Events**—Number of overrun events. • **CE Underrun Events**—Number of underrun events.	**extensive**

Sample Output

show interfaces (T1, IMA Link)

```
user@host> show interfaces t1-1/0/0
  IMA Link alarms    : None
  IMA Link defects   : LIF, LODS
  IMA Link state:
    Line     : Not synchronized
    Near end : Rx: Unusable, Tx: Usable
    Far  end : Rx: Unusable, Tx: Usable
  IMA link media:          Seconds        Count  State
    LIF                                       0  OK
    LODS                                      0  OK
    Err-ICP                                   0  OK
    IV                                        0  OK
    Rx-FC                                     0  OK
    Tx-FC                                     0  OK
    FE-Defects                                0
    FE-Rx-FC                                  0
    FE-Tx-FC                                  0
    Rx-ICP                                    0
    Rx-Stuff                                  0
    Tx-ICP                                   11
    Tx-Stuff                                  0
    Rx-SES                   0
    Rx-UAS                   0
    Rx-UUS                   1
    Tx-UUS                   0
    FE-Rx-SES                0
    FE-Rx-UAS                0
    FE-Rx-UUS                0
    FE-Tx-UUS                0
```

Copyright © 2018, Juniper Networks, Inc.

show interfaces (T1, PPP)

```
user@host> show interfaces t1-1/1/0
Physical interface: t1-1/1/0, Enabled, Physical link is Up
  Interface index: 149, SNMP ifIndex: 45
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: T1,
  Loopback: None, FCS: 16, Framing: ESF
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 0 (never), Output: 0 (never)
  LCP state: Opened
  NCP state: Opened
  CHAP state: Opened
  CoS queues     : 4 supported, 4 in use
  Last flapped   : 2005-12-05 08:43:06 PST (02:13:35 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 72 bps (0 pps)
  DS1   alarms   : None
  DS1   defects  : None

  Logical interface t1-1/1/0.0 (Index 66) (SNMP ifIndex 51)
    Flags: Hardware-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: 198.51.100.0/24, Local: 198.51.100.1, Broadcast:
198.51.100.255
```

show interfaces detail (T1, PPP)

```
user@host> show interfaces t1-1/1/0 detail
Physical interface: t1-1/1/0, Enabled, Physical link is Up
  Interface index: 149, SNMP ifIndex: 45, Generation: 32
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: T1,
  Loopback: None, FCS: 16, Framing: ESF
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 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 : 0 (last seen: never)
    Output: 0 (last sent: never)
  LCP state: Opened
  NCP state: Opened
  CHAP state: Opened
  CoS queues     : 4 supported, 4 in use
  Last flapped   : 2005-12-05 08:43:06 PST (02:13:52 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                  798                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                   42                    0 pps
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                    0                    0                    0
```


Copyright © 2018, Juniper Networks, Inc.

```
    1 expedited-fo                0                    0                    0

    2 assured-forw                0                    0                    0

    3 network-cont               40                   40                    0

  DS1   alarms   : None
  DS1   defects  : None
  DS1 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Induced Error rate: 10e-0, Algorithm: 2^15 - 1
  Logical interface t1-1/1/0.0 (Index 66) (SNMP ifIndex 51) (Generation 5)
    Flags: Hardware-Down Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 1500, Generation: 14, Route table: 0
      Flags: Protocol-Down
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 198.51.100.0/24, Local: 198.51.100.1, Broadcast:
198.51.100.255,
        Generation: 18
```

show interfaces extensive (T1 CRC Errors)

```
user@host> show interfaces t1-3/2/0:1:1 extensive
Physical interface: t1-3/2/0:1:1, Enabled, Physical link is Down
  Interface index: 179, SNMP ifIndex: 79, Generation: 180
  :
  :
  DS1   alarms   : AIS, LOF, CRC Major, CRC Minor
  DS1   defects  : AIS, LOF, CRC Major, CRC Minor
  T1  media:            Seconds        Count  State
    SEF                       1            1  OK
    BEE                       1            1  OK
    AIS                    1128            1  Defect Active
    LOF                    1128            1  Defect Active
    LOS                       0            0  OK
    YELLOW                    0            0  OK
    CRC Major               154            1  Defect Active
    CRC Minor               154            1  Defect Active
    BPV                       0            0
    EXZ                       0            0
    LCV                       0            0
    PCV                       0            0
    CS                        0            0
    CRC                     154        15400
    ...
```

show interfaces extensive (T1, PPP)

```
user@host> show interfaces t1-1/1/0 extensive
Physical interface: t1-1/1/0, Enabled, Physical link is Up
  Interface index: 149, SNMP ifIndex: 45, Generation: 32
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: T1,
  Loopback: None, FCS: 16, Framing: ESF
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 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 : 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     : 4 supported, 4 in use
Last flapped   : 2005-12-05 08:43:06 PST (02:13:54 ago)
Statistics last cleared: Never
Traffic statistics:
 Input  bytes  :                    0                  0 bps
 Output bytes  :                  817                 72 bps
 Input  packets:                    0                  0 pps
 Output packets:                   43                  0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  HS link CRC errors: 0, SRAM errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, Aged packets: 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                  42                   42                    0

DS1   alarms   : None
DS1   defects  : None
T1  media:            Seconds        Count  State
  SEF                       1            1  OK
  BEE                       0            0  OK
  AIS                       0            0  OK
  LOF                       1            1  OK
  LOS                       0            0  OK
  YELLOW                    1            1  OK
  BPV                       1            1
  EXZ                       1            1
  LCV                       1        65535
  PCV                       1         1023
  CS                        0            0
  LES                       1
  ES                        1
  SES                       1
  SEFS                      1
  BES                       0
  UAS                       0
HDLC configuration:
  Policing bucket: Disabled
  Shaping bucket : Disabled
  Giant threshold: 1514, Runt threshold: 3
  Timeslots      : All active
  Line encoding: B8ZS
  Buildout       : 0 to 132 feet
  Byte encoding: Nx64K, Data inversion: Disabled, Idle cycle flag: flags,
  Start end flag: shared
DS1 BERT configuration:
```

 Copyright © 2018, Juniper Networks, Inc.

```
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Induced Error rate: 10e-0, Algorithm: 2^15 - 1
  Packet Forwarding Engine configuration:
    Destination slot: 1, PLP byte: 1 (0x00)
  CoS information:
    CoS transmit queue       Bandwidth          Buffer   Priority    Limit
                           %          bps     %     usec
    0 best-effort         95      1459200    95         0        low    none
    3 network-control      5        76800     5         0        low    none

  Logical interface t1-1/1/0.0 (Index 66) (SNMP ifIndex 51) (Generation 5)
    Flags: Hardware-Down Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 1500, Generation: 14, Route table: 0
      Flags: Protocol-Down
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 198.51.100.0/24, Local: 198.51.100.1, Broadcast:
198.51.100.255,
        Generation: 18
```

show interfaces (E1, Frame Relay)

```
user@host> show interfaces e1-3/0/0
Physical interface: e1-3/0/0, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 37
  Link-level type: Frame-Relay, MTU: 1504, Clocking: Internal, Speed: E1,
  Loopback: None, FCS: 16, Framing: G704
  Device flags   : Present Running
  Interface flags: Link-Layer-Down Point-To-Point SNMP-Traps 16384
  Link flags     : Keepalives DTE
  ANSI LMI settings: n391dte 6, n392dte 3, n393dte 4, t391dte 10 seconds
  LMI: Input: 0 (never), Output: 11 (00:00:05 ago)
  DTE statistics:
    Enquiries sent                       : 10
    Full enquiries sent                  : 1
    Enquiry responses received           : 0
    Full enquiry responses received      : 0
  DCE statistics:
    Enquiries received                   : 0
    Full enquiries received              : 0
    Enquiry responses sent               : 0
    Full enquiry responses sent          : 0
  Common statistics:
    Unknown messages received            : 0
    Asynchronous updates received        : 0
    Out-of-sequence packets received     : 0
    Keepalive responses timedout         : 1
  CoS queues     : 8 supported
  Last flapped   : 2005-11-30 14:50:34 PST (4d 20:33 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  DS1   alarms   : None
  DS1   defects  : None
  Logical interface e1-3/0/0.0 (Index 72) (SNMP ifIndex 32)
    Flags: Device-Down Point-To-Point SNMP-Traps Encapsulation: FR-NLPID
  Input packets : 0
  Output packets: 0
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.1.3/24, Local: 10.1.3.1, Broadcast: 10.1.3.255
    DLCI 100
```

```
        Flags: Down, DCE-Unconfigured
        Total down time: 00:01:13 sec, Last down: 00:01:13 ago
        Input packets : 0
        Output packets: 0
    DLCI statistics:
      Active DLCI  :0  Inactive DLCI  :1
```

show interfaces detail (E1, Frame Relay)

```
user@host> show interfaces e1-3/0/0 detail
Physical interface: e1-3/0/0, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 37, Generation: 69
  Link-level type: Frame-Relay, MTU: 1504, Clocking: Internal, Speed: E1,
  Loopback: None, FCS: 16, Framing: G704
  Device flags   : Present Running
  Interface flags: Link-Layer-Down Point-To-Point SNMP-Traps 16384
  Link flags     : Keepalives DTE
  Hold-times     : Up 0 ms, Down 0 ms
  ANSI LMI settings: n391dte 6, n392dte 3, n393dte 4, t391dte 10 seconds
  LMI statistics:
    Input : 0 (last seen: never)
    Output: 12 (last sent 00:00:02 ago)
  DTE statistics:
    Enquiries sent                        : 10
    Full enquiries sent                   : 2
    Enquiry responses received            : 0
    Full enquiry responses received       : 0
  DCE statistics:
    Enquiries received                    : 0
    Full enquiries received               : 0
    Enquiry responses sent                : 0
    Full enquiry responses sent           : 0
  Common statistics:
    Unknown messages received             : 0
    Asynchronous updates received         : 0
    Out-of-sequence packets received      : 0
    Keepalive responses timedout          : 1
  CoS queues     : 8 supported
  Last flapped   : 2005-11-30 14:50:34 PST (4d 20:33 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                  225                   56 bps
   Input  packets:                    0                    0 pps
   Output packets:                   15                    0 pps
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 limited                        0                    0                    0

    1 expedited-fo                   0                    0                    0

    2 real-plus                      0                    0                    0

    3 network-cont                  15                   15                    0

  DS1   alarms   : None
  DS1   defects  : None
  DS1 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Induced Error rate: 10e-0, Algorithm: 2^15 - 1, 0.151, Pseudorandom (9)
  Logical interface e1-3/0/0.0 (Index 72) (SNMP ifIndex 32) (Generation 26)
```

 Copyright © 2018, Juniper Networks, Inc.

```
    Flags: Device-Down Point-To-Point SNMP-Traps Encapsulation: FR-NLPID
    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: 32, Route table: 0
      Flags: None
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.1.3/24, Local: 10.1.3.1, Broadcast: 10.1.3.255,
        Generation: 42
    DLCI 100
      Flags: Down, DCE-Unconfigured
      Total down time: 00:01:18 sec, Last down: 00:01:18 ago
      Traffic statistics:
       Input  bytes  :                   0
       Output bytes  :                   0
       Input  packets:                   0
       Output packets:                   0
  DLCI statistics:
    Active DLCI  :0  Inactive DLCI  :1
```

show interfaces extensive (E1, Frame Relay)

```
user@host> show interfaces e1-3/0/0 extensive
Physical interface: e1-3/0/0, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 37, Generation: 69
  Link-level type: Frame-Relay, MTU: 1504, Clocking: Internal, Speed: E1,
  Loopback: None, FCS: 16, Framing: G704
  Device flags   : Present Running
  Interface flags: Link-Layer-Down Point-To-Point SNMP-Traps 16384
  Link flags     : Keepalives DTE
  Hold-times     : Up 0 ms, Down 0 ms
  ANSI LMI settings: n391dte 6, n392dte 3, n393dte 4, t391dte 10 seconds
  LMI statistics:
    Input : 0 (last seen: never)
    Output: 12 (last sent 00:00:05 ago)
  DTE statistics:
    Enquiries sent                      : 10
    Full enquiries sent                 : 2
    Enquiry responses received          : 0
    Full enquiry responses received     : 0
  DCE statistics:
    Enquiries received                  : 0
    Full enquiries received             : 0
    Enquiry responses sent              : 0
    Full enquiry responses sent         : 0
  Common statistics:
    Unknown messages received           : 0
    Asynchronous updates received       : 0
    Out-of-sequence packets received    : 0
```

Copyright © 2018, Juniper Networks, Inc.

```
  Keepalive responses timedout         : 1
CoS queues     : 8 supported
Last flapped   : 2005-11-30 14:50:34 PST (4d 20:33 ago)
Statistics last cleared: Never
Traffic statistics:
 Input  bytes  :                   0                   0 bps
 Output bytes  :                 225                   0 bps
 Input  packets:                   0                   0 pps
 Output packets:                  15                   0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  HS link CRC errors: 0, SRAM errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 17, Errors: 0, Drops: 0, Aged packets: 0,
  MTU errors: 0, Resource errors: 0
Queue counters:       Queued packets  Transmitted packets      Dropped packets

  0 limited                        0                    0                    0

  1 expedited-fo                   0                    0                    0

  2 real-plus                      0                    0                    0

  3 network-cont                  15                   15                    0

DS1   alarms   : None
DS1   defects  : None
E1  media:            Seconds        Count  State
  SEF                       0            0  OK
  BEE                       5            5  OK
  AIS                       0            0  OK
  LOF                     245           15  OK
  LOS                     245            4  OK
  YELLOW                    0           11  OK
  BPV                       0            0
  EXZ                       9            9
  LCV                       0            0
  PCV                       0            0
  CS                        0            0
  FEBE                      0            0
  LES                       0
  ES                        0
  SES                       0
  SEFS                      0
  BES                       0
  UAS                     271
HDLC configuration:
  Policing bucket: Disabled
  Shaping bucket : Disabled
  Giant threshold: 1506, Runt threshold: 0
  Timeslots      : All active
  Line encoding: HDB3, Data inversion: Disabled, Idle cycle flag: flags,
  Start end flag: shared
DS1 BERT configuration:
  BERT time period: 10 seconds, Elapsed: 0 seconds
  Induced Error rate: 10e-0, Algorithm: 2^15 - 1, 0.151, Pseudorandom (9)
Packet Forwarding Engine configuration:
  Destination slot: 3, PLP byte: 1 (0x00)
CoS information:
  CoS transmit queue       Bandwidth          Buffer   Priority    Limit
```

 Copyright © 2018, Juniper Networks, Inc.

```
                     %          bps      %      usec
  0 limited          95      1945600     95        0        low      none
  3 network-control   5       102400      5        0        low      none
Logical interface e1-3/0/0.0 (Index 72) (SNMP ifIndex 32) (Generation 26)
  Flags: Device-Down Point-To-Point SNMP-Traps Encapsulation: FR-NLPID
  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: 32, Route table: 0
    Flags: None
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 10.1.3/24, Local: 10.1.3.1, Broadcast: 10.1.3.255,
      Generation: 42
  DLCI 100
    Flags: Down, DCE-Unconfigured
    Total down time: 00:01:21 sec, Last down: 00:01:21 ago
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
DLCI statistics:
  Active DLCI   :0  Inactive DLCI   :1
```

show interfaces (E1, IMA Link)

```
user@host> show interfaces e1-1/0/0
  IMA Link alarms    : None
  IMA Link defects   : LIF, LODS
  IMA Link state:
    Line      : Not synchronized
    Near end : Rx: Unusable, Tx: Usable
    Far  end : Rx: Unusable, Tx: Usable
  IMA link media:          Seconds        Count  State
    LIF                                       0
    LODS                                      0
    Err-ICP                                   0
    IV                                        0
    Rx-FC                                     0
    Tx-FC                                     0
    FE-Defects                                0
    FE-Rx-FC                                  0
    FE-Tx-FC                                  0
    Rx-ICP                                    0
    Rx-Stuff                                  0
    Tx-ICP                                   11
    Tx-Stuff                                  0
    Rx-SES                      0
    Rx-UAS                      0
```

Copyright © 2018, Juniper Networks, Inc.

```
    Rx-UUS                        1
    Tx-UUS                        0
    FE-Rx-SES                     0
    FE-Rx-UAS                     0
    FE-Rx-UUS                     0
    FE-Tx-UUS                     0
```

show interfaces extensive (T1, TDM-CCC-SATOP)

```
user@host> show interfaces t1-1/0/0:1:1 extensive
Physical interface: t1-1/0/0:1:1, Enabled, Physical link is Down
  Interface index: 153, SNMP ifIndex: 579, Generation: 817
  Link-level type: TDM-CCC-SATOP, MTU: 1504, Clocking: Internal, Speed: T1,
Loopback: None, FCS: 16, Framing: ESF,
  Parent: coc1-1/0/0:1 Interface index 152
  Device flags   : Present Running Down
  Interface flags: Hardware-Down Point-To-Point SNMP-Traps Internal: 0x0
  Link flags     : None
  Hold-times     : Up 0 ms, Down 0 ms
  CoS queues     : 8 supported, 8 maximum usable queues
  Last flapped   : 2012-10-28 02:12:40 PDT (22:32:13 ago)
  Statistics last cleared: 2012-10-29 00:44:52 PDT (00:00:01 ago)
  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
  DS1   alarms   : None
  DS1   defects  : None
  T1  media:            Seconds        Count  State
    SEF                       0            0   OK
    BEE                       0            0   OK
    AIS                       0            0   OK
    LOF                       0            0   OK
    LOS                       0            0   OK
    YELLOW                    0            0   OK
    CRC Major                 0            0   OK
    CRC Minor                 0            0   OK
    BPV                       0            0
    EXZ                       0            0
    LCV                       0            0
    PCV                       0            0
    CS                        0            0
    CRC                       0            0
    LES                       0
    ES                        0
    SES                       0
    SEFS                      0
    BES                       0
    UAS                       0
```


Copyright © 2018, Juniper Networks, Inc.

```
  SAToP configuration:
    Payload size: 192
    Idle pattern: 0xFF
    Octet aligned: Disabled
    Jitter buffer: packets: 8, latency: 7 ms, auto adjust: Disabled
    Excessive packet loss rate: sample period: 10000 ms, threshold: 30%
  DS1 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Induced Error rate: 0, Algorithm: 2^15 - 1, 0.151, Pseudorandom (9)
  SONET alarms   : None
  SONET defects  : AIS-V, RDI-V
  SONET vt:
    BIP-BIP2                     0            0
    REI-V                        0            0
    LOP-V                        0            0  OK
    AIS-V                        2            0  Defect Active
    RDI-V                        2            0  Defect Active
    UNEQ-V                       0            0  OK
    PLM-V                        0            0  OK
    ES-V                         0
    SES-V                        0
    UAS-V                        2
    ES-VFE                       0
    SES-VFE                      0
    UAS-VFE                      0
  Received SONET overhead:
V5      : 0x07
V5(cmp) : 0x02
  Transmitted SONET overhead:
V5      : 0x02
  Packet Forwarding Engine configuration:
    Destination slot: 1
  CoS information:
    Direction : Output
    CoS transmit queue              Bandwidth              Buffer Priority
Limit
                              %            bps     %           usec
    0 best-effort            95        1459200    95              0      low
none
    3 network-control         5          76800     5              0      low
none

  Logical interface t1-1/0/0:1:1.0 (Index 69) (SNMP ifIndex 580) (Generation 525)

    Flags: Device-Down Point-To-Point SNMP-Traps Encapsulation: TDM-CCC-SATOP
  CE info              Packets        Bytes  Count
    CE Tx                 1005        192960
    CE Rx                 1004        192768
    CE Rx Forwarded                       0
    CE Strayed               0
    CE Lost                  0
    CE Malformed             0
    CE Misinserted           0
    CE AIS dropped           0
    CE Dropped            1005        192960
    CE Overrun Events                        0
    CE Underrun Events                       0
    Protocol ccc, MTU: 1504, Generation: 814, Route table: 0
      Flags: Is-Primary
```

Copyright © 2018, Juniper Networks, Inc.

show interfaces extensive (DS, TDM-CCC-CESoPSN)

```
user@host> show interfaces ds-1/0/0:1:1:1 extensive
Physical interface: ds-1/0/0:1:1:1, Enabled, Physical link is Down
  Interface index: 154, SNMP ifIndex: 597, Generation: 819
  Link-level type: TDM-CCC-CESoPSN, MTU: 1504, Speed: 1536kbps, Loopback: None,
FCS: 16, Parent: ct1-1/0/0:1:1 Interface index 153
  Device flags   : Present Running Down
  Interface flags: Hardware-Down Point-To-Point SNMP-Traps Internal: 0x0
  Link flags     : None
  Hold-times     : Up 0 ms, Down 0 ms
  CoS queues     : 8 supported, 8 maximum usable queues
  Last flapped   : 2012-10-29 00:49:03 PDT (00:00:35 ago)
  Statistics last cleared: Never
  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
  CESoPSN configuration:
    Packetization latency: 1000 us
    Idle pattern: 0xFF
    Jitter buffer: packets: 8, latency: 8 ms, auto adjust: Disabled
    Excessive packet loss rate: sample period: 10000 ms, threshold: 30%
  DS0 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Induced Error rate: 0, Algorithm: 2^15 - 1, 0.151, Pseudorandom (9)
  Packet Forwarding Engine configuration:
    Destination slot: 1
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority
Limit
                              %            bps     %           usec
    0 best-effort            95        1459200    95              0       low
none
    3 network-control         5          76800     5              0       low
none

  Logical interface ds-1/0/0:1:1:1.0 (Index 69) (SNMP ifIndex 598) (Generation
549)
    Flags: Device-Down Point-To-Point SNMP-Traps Encapsulation: TDM-CCC-CESoPSN
  CE info               Packets         Bytes   Count
    CE Tx                     0             0
    CE Rx                 35712       6856704
    CE Rx Forwarded                         0
    CE Strayed                0
    CE Lost                   0
    CE Malformed              0
    CE Misinserted            0
```

 Copyright © 2018, Juniper Networks, Inc.

```
CE AIS dropped            0
CE Dropped                0            0
CE Overrun Events                         0
CE Underrun Events                        1
Protocol ccc, MTU: 1504, Generation: 857, Route table: 0
  Flags: Is-Primary
```

show interfaces (T3 or E3)

Syntax

show interfaces *interface-type*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display status information about the specified T3 or E3 interface.

Options

interface-type—On M Series and T Series routers, the T3 interface type is **t3-*fpc/pic/port***, whereas the E3 interface type is **e3-*fpc/pic/port***.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

List of Sample Output

show interfaces (T3, PPP) on page 1537
show interfaces detail (T3, PPP) on page 1538
show interfaces extensive (T3, PPP) on page 1539
show interfaces (E3, Frame Relay) on page 1540
show interfaces detail (E3, Frame Relay) on page 1541
show interfaces extensive (E3, Frame Relay) on page 1543

Output Fields

Table 76 on page 1528 lists the output fields for the **show interfaces** (T3 or E3) command. Output fields are listed in the approximate order in which they appear.

Table 76: T3 or E3 show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under “Common Output Fields Description” on page 1110.	All levels

 Copyright © 2018, Juniper Networks, Inc.

Table 76: T3 or E3 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
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**
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**.	All levels
Speed	Speed at which the interface is running.	All levels
Loopback	Whether loopback is enabled and the type of loopback (**local** or **remote**).	All levels
FCS	Frame check sequence on the interface (either **16** or **32**). The default is **16** bits.	All levels
Mode	(T3 only) Whether C-bit parity mode or M13 mode is enabled.	All levels
Long buildout	(T3 only) Buildout setting: less than 255 feet (68 meters) or greater than 255 feet and shorter than 450 feet (137 meters).	All levels
Framing	(E3 only) Physical layer framing format used on the link. It can be **G751** or **Unframed**. The default is **G751.**	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 1110.	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 1110.	All levels
Link flags	Information about the link. Possible values are described in the “Link Flags” section under “Common Output Fields Description” on page 1110.	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 keepalives. • **interval** ***seconds***—Time in seconds between successive keepalive requests. The range is **10** seconds through **32,767** seconds, with a default of **10** seconds. • **down-count** ***number***—Number of keepalive packets a destination must fail to receive before the network takes a link down. The range is **1** through **255**, with a default of **3**. • **up-count** ***number***—Number of keepalive packets a destination must receive to change a link's status from down to up. The range is **1** through **255**, with a default of **1**.	**detail extensive** none

Table 76: T3 or E3 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output

Keepalive statistics or **Keepalive**	(PPP and HDLC) Information about keepalive packets.
	• **Input**—Number of keepalive packets received by PPP.
	  ◦ **(last seen 00:00:00 ago)**—Time since the last keepalive packet was received, in the format ***hh*:*mm*:*ss***.
	• **Output**—Number of keepalive packets sent by PPP and how long ago the last keepalive packets were sent and received.
	  ◦ **(last seen 00:00:00 ago)**—Time since the last keepalive packet was sent, in the format ***hh*:*mm*:*ss***.
		detail extensive none

LMI settings	(Frame Relay) Local Management Interface (LMI) settings (ANSI or ITU). ANSI LMI settings is the default. The format is **LMI settings: *value*, *value*... *xx*** seconds, where ***value*** can be: • **n391dte**—DTE full status polling interval (1–255) • **n392dce**—DCE error threshold (1–10) • **n392dte**—DTE error threshold (1–10) • **n393dce**—DCE monitored event count (1–10) • **n393dte**—DTE monitored event count (1–10) • **t391dte**—DTE polling timer (5–30 seconds) • **t392dce**—DCE polling verification timer (5–30 seconds)	**detail extensive** none
LMI	(Frame Relay) LMI statistics: • **Input**—Number of packets coming in on the interface (*nn*) and how much time has passed since the last packet arrived. The format is **Input: *nn* (last seen *hh*:*mm*:*ss* ago)**. • **Output**—Number of packets sent out on the interface (*nn*) and how much time has passed since the last packet was sent. The format is **Output: *nn* (last sent *hh*:*mm*:*ss* ago)**.	**detail extensive** none
DTE statistics	(Frame Relay) Statistics about messages transmitted from the data terminal equipment (DTE) to the data communications equipment (DCE): • **Enquiries sent**—Number of link status enquiries sent from the DTE to the DCE. • **Full enquiries sent**—Number of full enquiries sent from the DTE to the DCE. • **Enquiry responses received**—Number of enquiry responses received by the DTE from the DCE. • **Full enquiry responses received**—Number of full enquiry responses sent from the DTE to the DCE.	**detail extensive** none
DCE statistics	(Frame Relay) Statistics about messages transmitted from the DCE to the DTE: • **Enquiries received**—Number of enquiries received by the DCE from the DTE. • **Full enquiries received**—Number of full enquiries received by the DCE from the DTE. • **Enquiry responses sent**—Number of enquiry responses sent from the DCE to the DTE. • **Full enquiry responses sent**—Number of full enquiry responses sent from the DCE to the DTE.	**detail extensive** none

Copyright © 2018, Juniper Networks, Inc.

Table 76: T3 or E3 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Common statistics	(Frame Relay) Statistics about messages sent between the DTE and the DCE: • **Unknown messages received**—Number of received packets that do not fall into any category. • **Asynchronous updates received**—Number of link status peer changes received. • **Out-of-sequence packets received**—Number of packets for which the sequence of the packets received is different from the expected sequence. • **Keepalive responses timedout**—Number of keepalive responses that timed out when no LMI packet was reported for **n392dte** or **n393dce** intervals. (See **LMI settings**.)	**detail extensive** none
Nonmatching DCE-end DLCIs	(Frame Relay. Displayed only from the DTE.) Number of DLCIs configured from the DCE.	**detail extensive** none
LCP state	(PPP) Link Control Protocol state. • **Conf-ack-received**—Acknowledgement was received. • **Conf-ack-sent**—Acknowledgement was sent. • **Conf-req-sent**—Request was sent. • **Down**—LCP negotiation is incomplete (not yet completed or has failed). • **Opened**—LCP negotiation is successful.	**detail extensive** none
NCP state	(PPP) Network Control Protocol state. • **Conf-ack-received**—Acknowledgement was received. • **Conf-ack-sent**—Acknowledgement was sent. • **Conf-req-sent**—Request was sent. • **Down**—NCP negotiation is incomplete (not yet completed or has failed). • **Opened**—NCP negotiation is successful.	**detail extensive** none
CHAP state	(PPP) Displays the state of the Challenge Handshake Authentication Protocol (CHAP) during its transaction. • **Chap-Resp-received**—Response received for the challenge sent, but CHAP not yet moved into the Success state. (Most likely with RADIUS authentication.) • **Chap-Resp-sent**—Response sent for the challenge received. • **Chap-Chal-sent**—Challenge sent. • **Chap-Chal-received**—Challenge received but response not yet sent. • **Down**—CHAP authentication is incomplete (not yet completed or has failed). • **Not-configured**—CHAP is not configured on the interface. • **Opened**—CHAP authentication was successful.	**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*** **(*year-month-day hour:minute:second* ago)** . For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.	**detail extensive** none
CoS queues	Number of CoS queues configured.	**detail extensive** none

Table 76: T3 or E3 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Input rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface • **Output packets**—Number of packets received on the interface.	**detail extensive**
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—(T3 only) Number of frames received that are smaller than the runt threshold. • **Giants**—(T3 only) Number of frames received that are larger than the giant threshold. • **Bucket Drops**—Drops resulting from the traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is **off**. • **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. • **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. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • **SRAM errors**—Number of hardware errors that occurred in the static RAM (SRAM) on the PIC or PIM. If the value of this field increments, the PIC or PIM is malfunctioning. • **Resource errors**—Sum of transmit drops.	**extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 76: T3 or E3 show interfaces Output Fields (continued)

<table>
<tr><th>Field Name</th><th>Field Description</th><th>Level of Output</th></tr>
<tr><td>Output errors</td><td>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:
Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments (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.
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.
MTU errors—Number of packets whose size exceeded the MTU of the interface.
Resource errors—Sum of transmit drops.</td><td>extensive</td></tr>
<tr><td>Queue counters</td><td>CoS queue number and its associated user-configured forwarding class name.
Queued packets—Number of queued packets.
Transmitted packets—Number of transmitted packets.
Dropped packets—Number of packets dropped by the ASIC's RED mechanism.</td><td>detail extensive</td></tr>
<tr><td>Active alarms

Active defects</td><td>E3 media-specific defects that can render the interface unable to pass 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.
AIS—Alarm indication signal
EXZ—Excessive zeros
FERF—Far-end receive failures
IDLE—Idle code detected
LCD—Loss of cell delineation
LCV—Line code violation
LOF—Loss of frame
LOS—Loss of signal
PLL—Phase-locked loop
YLW—Remote defect indication</td><td>detail extensive none</td></tr>
</table>

Table 76: T3 or E3 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
DS3 media or **E3 media**	Counts of DS3 (T3) or E3 media-specific errors. • **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. The DS3 or E3 media-specific error types can be: ◦ **PLL Lock**—Phase-locked loop out of lock ◦ **Reframing**—Frame alignment recovery time ◦ **AIS**—Alarm indication signal ◦ **LOF**—Loss of frame ◦ **LOS**—Loss of signal ◦ **IDLE**—Idle code detected ◦ **YELLOW**—Errors at the remote site receiver ◦ **BPV**—Bipolar violation ◦ **EXZ**—Excessive zeros ◦ **LCV**—Line code violation ◦ **PCV**—(DS3 only) Pulse code violation ◦ **CCV**—(DS3 only) C-bit coding violation ◦ **FEBE**—(DS3 only) Far-end block error ◦ **LES**—Line error seconds ◦ **PES**—(DS3 only) P-bit errored seconds ◦ **PSES**—(DS3 only) P-bit errored seconds (section) ◦ **CES**—(DS3 only) C-bit errored seconds ◦ **CSES**—(DS3 only) C-bit severely errored seconds ◦ **SEFS**—Severely errored framing seconds ◦ **UAS**—Unavailable seconds	**extensive**
HDLC configuration	Information about the HDLC configuration. • **Policing bucket**—Configured state of the receiving policer. • **Shaping bucket**—Configured state of the transmitting shaper. • **Giant threshold**—Giant threshold programmed into the hardware. • **Runt threshold**—Runt threshold programmed into the hardware. • **Idle cycle flag**—Idle cycle flags. • **Start end flag**—Start and end flag.	**extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 76: T3 or E3 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
DSU configuration	Information about the DSU configuration. The last three lines (**Bit count**, **Error bit count**, and **LOS information**) are displayed only if a BERT has ever been run on the interface. • **Compatibility mod**—CSU/DSU compatibility mode: **None**, **Larscom**, **Kentrox**, or **Digital-Link**. • **Scrambling**—Payload scrambling: **Enabled** or **Disabled**. • **Subrate**—Configured subrate setting. Applies only when **Digital-Link** compatibility mode is used. The subrate can be **Disabled** or display units in Kbps. • **FEAC loopbac**—(T3) Whether a far-end alarm and control (FEAC) loopback is **Active** or **Inactive**. This feature is used to send alarm or status information from the far-end terminal back to the near-end terminal and to initiate T3 loopbacks at the far-end terminal from the near-end terminal. • **Response**—Whether the FEAC signal is **Enabled** or **Disabled**. • **Count**—Number of FEAC loopbacks.	**extensive**
DS3 (or **E3**) **BERT configuration**	BERT (bit error rate test) checks the quality of the line. This output appears only when a BERT is run on the interface. • **BERT time period**—Configured total time period that the BERT is to run. • **Elapsed**—Actual time elapsed since the start of the BERT (in seconds). • **Induced error rate**—Configured rate at which the bit errors are induced in the BERT pattern. • **Algorithm**—Type of algorithm selected for the BERT.	**detail extensive** none
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number. • **PLP byte**—Packet Level Protocol byte.	**extensive**
CoS information	Information about the CoS queue for the physical interface. • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**

Logical Interface

Table 76: T3 or E3 show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Logical interface	Name of the logical interface.	**detail extensive** none
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.	**extensive**
Flags	Information about the interface. Possible values are described in the "Interface Flags" section under "Common Output Fields Description" on page 1110.	**detail extensive** none
Encapsulation	Encapsulation on the logical interface.	**detail extensive** none
Input packets	Number of packets received on the logical interface.	None specified
Output packets	Number of packets transmitted on the logical interface.	None specified
Traffic statistics	(Frame Relay) Number and rate of bytes and packets received and transmitted on the logical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Local statistics	(Frame Relay) Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	**detail extensive**
Transit statistics	(Frame Relay) Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. This counter normally stabilizes in less than 1 second.	**detail extensive**
Protocol	Protocol family configured on the logical interface, such as **iso**, **inet6**, **mlfr**, or **mpls**.	**detail extensive** none
Multilink bundle	(Multilink) Interface name for the multilink bundle.	**detail extensive** none
MTU	MTU size on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route table	Routing table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive**
Flags	Information about the protocol family flags. Possible values are described in the "Family Flags" section under "Common Output Fields Description" on page 1110.	**detail extensive** none

 Copyright © 2018, Juniper Networks, Inc.

Table 76: T3 or E3 show interfaces 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 1110.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive** none
DLCI	(Frame Relay) DLCI number of the logical interface. The following DLCI information is displayed: **Flags**, **Total down time**, **Last down**, and **Traffic statistics** (or **Input packets**, **Output packets**). **Flags** is one or more of the following: • **Active**—Set when the link is active and the DTE and DCE are exchanging information. • **Down**—Set when the link is active, but no information is received from the DCE. • **DCE Unconfigured**—Set when the corresponding DLCI in the DCE is not configured. • **Configured**—Set when the corresponding DLCI in the DCE is configured. • **DCE-configured**—Displayed when the command is issued from the DTE.	**detail extensive** none
DLCI statistics	(Frame Relay) Data-link connection identifier (DLCI) statistics. • **Active DLCI**—Number of active DLCIs. • **Inactive DLCI**—Number of inactive DLCIs.	**detail extensive** none

Sample Output

show interfaces (T3, PPP)

```
user@host> show interfaces t3-0/2/0
Physical interface: t3-0/2/0, Enabled, Physical link is Up
  Interface index: 139, SNMP ifIndex: 35
  Link-level type: PPP, MTU: 4474, Clocking: Internal, Speed: T3,
  Loopback: None, FCS: 16, Mode: C/Bit parity,
  Long buildout: Shorter than 255 feet
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Link flags     : Keepalives
  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     : 4 supported, 4 in use
  Last flapped   : 2005-12-05 08:43:06 PST (02:18:40 ago)
```

Copyright © 2018, Juniper Networks, Inc.

```
  Input rate     : 0 bps (0 pps)
  Output rate    : 72 bps (0 pps)
  Active alarms  : None
  Active defects : None
  DS3 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Algorithm: 2^15 - 1, 0.151, Pseudorandom (9), Induced error rate: 10e-0

Logical interface t3-0/2/0.0 (Index 66) (SNMP ifIndex 54)
    Flags: Hardware-Down Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 4470
      Flags: Protocol-Down
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 198.51.100.0/24, Local: 198.51.100.1, Broadcast:
198.51.100.255
```

show interfaces detail (T3, PPP)

```
user@host> show interfaces t3-0/2/0 detail
Physical interface: t3-0/2/0, Enabled, Physical link is Up
  Interface index: 139, SNMP ifIndex: 35, Generation: 22
  Link-level type: PPP, MTU: 4474, Clocking: Internal, Speed: T3,
  Loopback: None, FCS: 16, Mode: C/Bit parity,
  Long buildout: Shorter than 255 feet
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 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 : 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     : 4 supported, 4 in use
  Last flapped   : 2005-12-05 08:43:06 PST (02:18:45 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                  152                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    8                    0 pps
  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                   6                    6                    0

  Active alarms  : None
  Active defects : None
  DS3 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Algorithm: 2^15 - 1, 0.151, Pseudorandom (9), Induced error rate: 10e-0

Logical interface t3-0/2/0.0 (Index 66) (SNMP ifIndex 54) (Generation 8)
```


Copyright © 2018, Juniper Networks, Inc.

```
    Flags: Hardware-Down Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 4470, Generation: 17, Route table: 0
      Flags: Protocol-Down
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 198.51.100.0/24, Local: 198.51.100.1, Broadcast:
198.51.100.255,
        Generation: 24
```

show interfaces extensive (T3, PPP)

```
user@host> show interfaces t3-0/2/0 extensive
Physical interface: t3-0/2/0, Enabled, Physical link is Up
  Interface index: 139, SNMP ifIndex: 35, Generation: 22
  Link-level type: PPP, MTU: 4474, Clocking: Internal, Speed: T3,
  Loopback: None, FCS: 16, Mode: C/Bit parity,
  Long buildout: Shorter than 255 feet
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 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 : 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     : 4 supported, 4 in use
  Last flapped   : 2005-12-05 08:43:06 PST (02:18:47 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                  171                   72 bps
   Input  packets:                    0                    0 pps
   Output packets:                    9                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Bucket drops: 0, Policed discards: 0, L3 incompletes: 0,
    L2 channel errors: 0, L2 mismatch timeouts: 0, HS link CRC errors: 0,
    SRAM errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, Aged packets: 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                  7                    7                    0

  Active alarms  : None
  Active defects : None
  DS3 media:            Seconds        Count  State
    PLL Lock                  0            0  OK
    Reframing                 0            0  OK
    AIS                       0            0  OK
```

```
    LOF                        0            0  OK
    LOS                        0            0  OK
    IDLE                       0            0  OK
    YELLOW                     0            0  OK
    BPV                        0            0
    EXZ                        0            0
    LCV                        1            4
    PCV                        0            0
    CCV                        0            0
    FEBE                       1           11
    LES                        1
    PES                        0
    PSES                       0
    CES                        0
    CSES                       0
    SEFS                       0
    UAS                        0
  HDLC configuration:
    Policing bucket: Disabled
    Shaping bucket : Disabled
    Giant threshold: 4484, Runt threshold: 3
    Idle cycle flag: flags, Start end flag: shared
  DSU configuration:
    Compatibility mode: None, Scrambling: Disabled, Subrate: Disabled
    FEAC loopback: Inactive, Response: Disabled, Count: 0
  DS3 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Algorithm: 2^15 - 1, 0.151, Pseudorandom (9), Induced error rate: 10e-0
  Packet Forwarding Engine configuration:
    Destination slot: 0, PLP byte: 1 (0x00)
  CoS information:
    CoS transmit queue       Bandwidth          Buffer   Priority    Limit
                         %          bps     %     usec
    0 best-effort       95    42499200     95        0         low    none
    3 network-control    5     2236800      5        0         low    none

Logical interface t3-0/2/0.0 (Index 66) (SNMP ifIndex 54) (Generation 8)
    Flags: Hardware-Down Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 4470, Generation: 17, Route table: 0
      Flags: Protocol-Down
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 198.51.100.0/24, Local: 198.51.100.1, Broadcast:
198.51.100.255,
        Generation: 24
```

show interfaces (E3, Frame Relay)

```
user@host> show interfaces e3-1/2/0
Physical interface: e3-1/2/0, Enabled, Physical link is Up
  Interface index: 153, SNMP ifIndex: 49
  Link-level type: Frame-Relay, MTU: 4474, Clocking: Internal, Speed: E3,
  Loopback: None, FCS: 16, Framing: G751
  Device flags   : Present Running
  Interface flags: Link-Layer-Down Point-To-Point SNMP-Traps Internal: 0x4000
  Link flags     : Keepalives DTE
  ANSI LMI settings: n391dte 6, n392dte 3, n393dte 4, t391dte 10 seconds
  LMI: Input: 0 (never), Output: 4 (00:00:06 ago)
  DTE statistics:
    Enquiries sent                     : 4
    Full enquiries sent                : 0
    Enquiry responses received         : 0
```

 Copyright © 2018, Juniper Networks, Inc.

```
    Full enquiry responses received      : 0
  DCE statistics:
    Enquiries received                   : 0
    Full enquiries received              : 0
    Enquiry responses sent               : 0
    Full enquiry responses sent          : 0
  Common statistics:
    Unknown messages received            : 0
    Asynchronous updates received        : 0
    Out-of-sequence packets received     : 0
    Keepalive responses timedout         : 1
  CoS queues     : 4 supported, 4 in use
  Last flapped   : 2005-12-05 08:46:14 PST (02:27:17 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None

  Logical interface e3-1/2/0.0 (Index 66) (SNMP ifIndex 57)
    Flags: Device-Down Point-To-Point SNMP-Traps Encapsulation: FR-NLPID
  Input packets : 0
  Output packets: 0
    Protocol inet, MTU: 4470
      Flags: None
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 198.51.100.0/24, Local: 198.51.100.1, Broadcast:
198.51.100.255
    DLCI 100
      Flags: Down, DCE-Unconfigured
      Total down time: 00:00:06 sec, Last down: 00:00:06 ago
      Input packets : 0
      Output packets: 0
  DLCI statistics:
    Active DLCI   :0  Inactive DLCI   :1
```

show interfaces detail (E3, Frame Relay)

```
user@host> show interfaces e3-1/2/0 detail
Physical interface: e3-1/2/0, Enabled, Physical link is Up
  Interface index: 153, SNMP ifIndex: 49, Generation: 36
  Link-level type: Frame-Relay, MTU: 4474, Clocking: Internal, Speed: E3,
  Loopback: None, FCS: 16, Framing: G751
  Device flags   : Present Running
  Interface flags: Link-Layer-Down Point-To-Point SNMP-Traps Internal: 0x4000
  Link flags     : Keepalives DTE
  Hold-times     : Up 0 ms, Down 0 ms
  ANSI LMI settings: n391dte 6, n392dte 3, n393dte 4, t391dte 10 seconds
  LMI statistics:
    Input : 0 (last seen: never)
    Output: 5 (last sent 00:00:07 ago)
  DTE statistics:
    Enquiries sent                       : 5
    Full enquiries sent                  : 0
    Enquiry responses received           : 0
    Full enquiry responses received      : 0
  DCE statistics:
    Enquiries received                   : 0
    Full enquiries received              : 0
    Enquiry responses sent               : 0
    Full enquiry responses sent          : 0
  Common statistics:
```

```
    Unknown messages received           : 0
    Asynchronous updates received       : 0
    Out-of-sequence packets received    : 0
    Keepalive responses timedout        : 1
  CoS queues     : 4 supported, 4 in use
  Last flapped   : 2005-12-05 08:46:14 PST (02:27:27 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                    0 bps
   Output bytes  :                 806                    0 bps
   Input  packets:                   0                    0 pps
   Output packets:                  44                    0 pps
  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                43                   43                    0

  Active alarms  : None
  Active defects : None

  Logical interface e3-1/2/0.0 (Index 66) (SNMP ifIndex 57) (Generation 15)
    Flags: Device-Down Point-To-Point SNMP-Traps Encapsulation: FR-NLPID
    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: 4470, Generation: 24, Route table: 0
      Flags: None
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 198.51.100.0/24, Local: 198.51.100.1, Broadcast:
198.51.100.255,
        Generation: 38
    DLCI 100
      Flags: Down, DCE-Unconfigured
      Total down time: 00:00:16 sec, Last down: 00:00:16 ago
      Traffic statistics:
       Input  bytes  :                   0
       Output bytes  :                   0
       Input  packets:                   0
       Output packets:                   0
  DLCI statistics:
    Active DLCI  :0  Inactive DLCI  :1
```

 Copyright © 2018, Juniper Networks, Inc.

show interfaces extensive (E3, Frame Relay)

```
user@host> show interfaces e3-1/2/0 extensive
Physical interface: e3-1/2/0, Enabled, Physical link is Up
  Interface index: 153, SNMP ifIndex: 49, Generation: 36
  Link-level type: Frame-Relay, MTU: 4474, Clocking: Internal, Speed: E3,
  Loopback: None, FCS: 16, Framing: G751
  Device flags   : Present Running
  Interface flags: Link-Layer-Down Point-To-Point SNMP-Traps Internal: 0x4000
  Link flags     : Keepalives DTE
  Hold-times     : Up 0 ms, Down 0 ms
  ANSI LMI settings: n391dte 6, n392dte 3, n393dte 4, t391dte 10 seconds
  LMI statistics:
    Input : 0 (last seen: never)
    Output: 6 (last sent 00:00:02 ago)
  DTE statistics:
    Enquiries sent                       : 5
    Full enquiries sent                  : 1
    Enquiry responses received           : 0
    Full enquiry responses received      : 0
  DCE statistics:
    Enquiries received                   : 0
    Full enquiries received              : 0
    Enquiry responses sent               : 0
    Full enquiry responses sent          : 0
  Common statistics:
    Unknown messages received            : 0
    Asynchronous updates received        : 0
    Out-of-sequence packets received     : 0
    Keepalive responses timedout         : 1
  CoS queues     : 4 supported, 4 in use
  Last flapped   : 2005-12-05 08:46:14 PST (02:27:30 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                    0 bps
   Output bytes  :                 821                   56 bps
   Input  packets:                   0                    0 pps
   Output packets:                  45                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 21118, Bucket drops: 0,
    Policed discards: 0, L3 incompletes: 0, L2 channel errors: 0,
    L2 mismatch timeouts: 0, HS link CRC errors: 0, SRAM errors: 0,
    Resource errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, Aged packets: 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                  44                   44                    0

  Active alarms  : None
  Active defects : None
  E3  media:            Seconds        Count  State
    PLL Lock                  0            0  OK
```

```
    Reframing              187           1  OK
    AIS                      0           0  OK
    LOF                    187           1  OK
    LOS                    187           1  OK
    IDLE                     0           0  OK
    YELLOW                   0           0  OK
    BPV                      0           0
    EXZ                      0           0
    LCV                    188    12303167
    LES                    188
    SEFS                   187
    UAS                    195
  DSU configuration:
    Compatibility mode: None, Scrambling: Disabled
  E3 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Algorithm: 2^15 - 1, O.151, Pseudorandom (9), Induced Error rate: 10e-0
  Packet Forwarding Engine configuration:
    Destination slot: 1, PLP byte: 1 (0x00)
  CoS information:
    CoS transmit queue       Bandwidth          Buffer   Priority    Limit
                        %          bps     %     usec
    0 best-effort      95     32649600    95        0         low     none
    3 network-control   5      1718400     5        0         low     none

  Logical interface e3-1/2/0.0 (Index 66) (SNMP ifIndex 57) (Generation 15)
    Flags: Device-Down Point-To-Point SNMP-Traps Encapsulation: FR-NLPID
    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: 4470, Generation: 24, Route table: 0
      Flags: None
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 198.51.100.0/24, Local: 198.51.100.1, Broadcast:
198.51.100.255,
        Generation: 38
    DLCI 100
      Flags: Down, DCE-Unconfigured
      Total down time: 00:00:19 sec, Last down: 00:00:19 ago
      Traffic statistics:
       Input  bytes  :                    0
       Output bytes  :                    0
       Input  packets:                    0
       Output packets:                    0
  DLCI statistics:
    Active DLCI  :0  Inactive DLCI  :1
```

 Copyright © 2018, Juniper Networks, Inc.

show interfaces demux0 (Demux Interfaces)

Syntax

show interfaces demux0.*logical-interface-number*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced in Junos OS Release 9.0.

Description

(MX Series and M Series routers only) Display status information about the specified demux interface.

Options

none—Display standard information about the specified demux interface.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index ***snmp-index***—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Related Documentation

- *Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*

List of Sample Output

show interfaces demux0 (Demux) on page 1551
show interfaces demux0 (PPPoE over Aggregated Ethernet) on page 1552
show interfaces demux0 extensive (Targeted Distribution for Aggregated Ethernet Links) on page 1553
show interfaces demux0 (ACI Interface Set Configured) on page 1553

Output Fields

Table 77 on page 1545 lists the output fields for the **show interfaces** demux0 (Demux Interfaces) command. Output fields are listed in the approximate order in which they appear.

Table 77: show interfaces demux0 (Demux Interfaces) Output Fields

Field Name	Field Description	Level of Output
Physical Interface		

Table 77: show interfaces demux0 (Demux Interfaces) Output Fields (continued)

Field Name	Field Description	Level of Output
Physical interface	Name of the physical interface.	**brief detail extensive** none
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**brief detail extensive** none
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under “Common Output Fields Description” on page 1110.	**brief detail extensive** none
Physical link	Status of the physical link (**Up** or **Down**).	**detail extensive** none
Admin	Administrative state of the interface (**Up** or **Down**).	**terse**
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**detail extensive** none
Link	Status of the physical link (**Up** or **Down**).	**terse**
Targeting summary	Status of aggregated Ethernet links that are configured with targeted distribution (**primary** or **backup**)	**extensive**
Bandwidth	Bandwidth allocated to the aggregated Ethernet links that are configured with targeted distribution.	**extensive**
Proto	Protocol family configured on the interface.	**terse**
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Type	Type of interface. **Software-Pseudo** indicates a standard software interface with no associated hardware device.	**brief detail extensive** none
Link-level type	Encapsulation being used on the physical interface.	**brief detail extensive**
MTU	Maximum transmission unit size on the physical interface.	**brief detail extensive**
Clocking	Reference clock source: **Internal** (1) or **External** (2).	**brief detail extensive**
Speed	Speed at which the interface is running.	**brief detail extensive**
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under “Common Output Fields Description” on page 1110.	**brief detail extensive** none
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 1110.	**brief detail extensive** none
Link type	Data transmission type.	**detail extensive** none

Copyright © 2018, Juniper Networks, Inc.

Table 77: show interfaces demux0 (Demux Interfaces) Output Fields (continued)

Field Name	Field Description	Level of Output
Link flags	Information about the link. Possible values are described in the "Link Flags" section under "Common Output Fields Description" on page 1110.	**detail extensive** none
Physical info	Information about the physical interface.	**detail extensive**
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive**
Hardware address	Hardware MAC address.	**detail extensive**
Alternate link address	Backup address of the link.	**detail extensive**
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour*:*minute*:*second*:*timezone*** **(*hour*:*minute*:*second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).**	**detail extensive** none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface. • **IPv6 transit statistics**—Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled. NOTE: These fields include dropped traffic and exception traffic, as those fields are not separately defined. ◦ **Input bytes**—Number of bytes received on the interface ◦ **Output bytes**—Number of bytes transmitted on the interface. ◦ **Input packets**—Number of packets received on the interface. ◦ **Output packets**—Number of packets transmitted on the interface.	**detail extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 77: show interfaces demux0 (Demux Interfaces) Output Fields (continued)

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface whose definitions are as follows: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Giants**—Number of frames received that are larger than the giant packet threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • **Resource errors**—Sum of transmit drops.	**extensive**
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	none
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**
Output Rate	Output rate in bps and pps.	none
Logical Interface		
Logical interface	Name of the logical interface.	**brief detail extensive** none
Index	Index number of the logical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP interface index number for the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail**
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under “Common Output Fields Description” on page 1110.	**brief detail extensive** none

 Copyright © 2018, Juniper Networks, Inc.

Table 77: show interfaces demux0 (Demux Interfaces) Output Fields (continued)

Field Name	Field Description	Level of Output
Encapsulation	Encapsulation on the logical interface.	**brief extensive** none
ACI VLAN: Dynamic Profile	Name of the dynamic profile that defines the agent circuit identifier (ACI) interface set. If configured, the ACI interface set enables the underlying demux interface to create dynamic VLAN subscriber interfaces based on ACI information.	**brief detail extensive** none
Demux	Specific IP demultiplexing (demux) values: • **Underlying interface**—The underlying interface that the demux interface uses. • **Index**—Index number of the logical interface. • **Family**—Protocol family configured on the logical interface. • **Source prefixes, total**—Total number of source prefixes for the underlying interface. • **Destination prefixes, total**—Total number of destination prefixes for the underlying interface. • **Prefix**—**inet** family prefix.	**detail extensive** none
protocol-family	Protocol family configured on the logical interface.	**brief**

Traffic statistics	Number and rate of bytes and packets received and transmitted on the specified interface set.
	• **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface set.
	• **Input packets, Output packets**—Number of packets received and transmitted on the interface set.
	• **IPv6 transit statistics**—Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.
	NOTE: The packet and byte counts in these fields include traffic that is dropped and does not leave the router.
	  ◦ **Input bytes**—Number of bytes received on the interface.
	  ◦ **Output bytes**—Number of bytes transmitted on the interface.
	  ◦ **Input packets**—Number of packets received on the interface.
	  ◦ **Output packets**—Number of packets transmitted on the interface.
		detail extensive

Local statistics	Number of transit bytes and packets received and transmitted on the local interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 77: show interfaces demux0 (Demux Interfaces) Output Fields (continued)

Field Name	Field Description	Level of Output
Transit statistics	Number and rate of bytes and packets transiting the switch. NOTE: The packet and byte counts in these fields include traffic that is dropped and does not leave the router. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
IPv6 Transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled. NOTE: The packet and byte counts in these fields include traffic that is dropped and does not leave the router. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Input packets	Number of packets received on the interface.	none
Output packets	Number of packets transmitted on the interface.	none
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under “Common Output Fields Description” on page 1110.	**detail extensive** none
MTU	Maximum transmission unit size on the logical interface.	**detail extensive** none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route table	Route table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive**
Flags	Information about protocol family flags. Possible values are described in the “Family Flags” section under “Common Output Fields Description” on page 1110.	**detail extensive** none
Mac-Validate Failures	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	**detail extensive** none
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under “Common Output Fields Description” on page 1110.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive statistics** none

 Copyright © 2018, Juniper Networks, Inc.

Table 77: show interfaces demux0 (Demux Interfaces) Output Fields (continued)

Field Name	Field Description	Level of Output
Local	IP address of the logical interface.	**detail extensive terse** none
Remote	IP address of the remote interface.	**terse**
Broadcast	Broadcast address of the logical interlace.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Link	Name of the physical interfaces for member links in an aggregated Ethernet bundle for a PPPoE over aggregated Ethernet configuration. PPPoE traffic goes out on these interfaces.	**detail extensive** none
Dynamic-profile	Name of the PPPoE dynamic profile assigned to the underlying interface.	**detail extensive** none
Service Name Table	Name of the PPPoE service name table assigned to the PPPoE underlying interface.	**detail extensive** none
Max Sessions	Maximum number of dynamic PPPoE logical interfaces that the router can activate on the underlying interface.	**detail extensive** none
Duplicate Protection	State of duplicate protection: **On** or **Off**. Duplicate protection prevents the activation of another dynamic PPPoE logical interface on the same underlying interface when a dynamic PPPoE logical interface for a client with the same MAC address is already active on that interface.	**detail extensive** none
Direct Connect	State of the configuration to ignore DSL Forum VSAs: **On** or **Off**. When configured, the router ignores any of these VSAs received from a directly connected CPE device on the interface.	**detail extensive** none
AC Name	Name of the access concentrator.	**detail extensive** none

Sample Output

show interfaces demux0 (Demux)

```
user@host> show interfaces demux0
Physical interface: demux0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 79, Generation: 129
  Type: Software-Pseudo, Link-level type: Unspecified, MTU: 9192, Clocking: 1,
  Speed: Unspecified
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
```

Copyright © 2018, Juniper Networks, Inc.

```
  Traffic statistics:
   Input  bytes  :               0                    0 bps
   Output bytes  :               0                    0 bps
   Input  packets:               0                    0 pps
   Output packets:               0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :              0
    Output bytes  :              0
    Input  packets:              0
    Output packets:              0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0

Logical interface demux0.0 (Index 87) (SNMP ifIndex 84) (Generation 312)
    Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
    Demux:
     Underlying interface: ge-2/0/1.0 (Index 74)
   Family Inet Source prefixes, total 1
   Prefix: 203.0.113/24
    Traffic statistics:
     Input  bytes  :               0
     Output bytes  :            1554
     Input  packets:               0
     Output packets:              37
     IPv6 transit statistics:
      Input  bytes  :              0
      Output bytes  :              0
      Input  packets:              0
      Output packets:              0
    Local statistics:
     Input  bytes  :               0
     Output bytes  :            1554
     Input  packets:               0
     Output packets:              37
    Transit statistics:
     Input  bytes  :               0                    0 bps
     Output bytes  :               0                    0 bps
     Input  packets:               0                    0 pps
     Output packets:               0                    0 pps
     IPv6 transit statistics:
      Input  bytes  :              0
      Output bytes  :              0
      Input  packets:              0
      Output packets:              0
    Protocol inet, MTU: 1500, Generation: 395, Route table: 0
      Flags: Is-Primary, Mac-Validate-Strict
      Mac-Validate Failures: Packets: 0, Bytes: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 203.0.113/24, Local: 203.0.113.13, Broadcast: 203.0.113.255,

        Generation: 434
```

show interfaces demux0 (PPPoE over Aggregated Ethernet)

```
user@host> show interfaces demux0.100
Logical interface demux0.100 (Index 76) (SNMP ifIndex 61160)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ]
```


Copyright © 2018, Juniper Networks, Inc.

```
    Encapsulation: ENET2
    Demux:
      Underlying interface: ae0 (Index 199)
    Link:
      ge-1/0/0
      ge-1/1/0
    Input packets : 0
    Output packets: 0
    Protocol pppoe
      Dynamic Profile: pppoe-profile,
      Service Name Table: service-table1,
      Max Sessions: 100, Duplicate Protection: On,
      Direct Connect: Off,
      AC Name: pppoe-server-1
```

show interfaces demux0 extensive (Targeted Distribution for Aggregated Ethernet Links)

user@host> **show interfaces demux0.1073741824 extensive**

```
  Logical interface demux0.1073741824 (Index 75) (SNMP ifIndex 558) (Generation
346)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ]  Encapsulation: ENET2
    Demux:
      Underlying interface: ae0 (Index 201)
    Link:
      ge-1/0/0
      ge-1/1/0
      ge-2/0/7
      ge-2/0/8
    Targeting summary:
      ge-1/1/0, primary, Physical link is Up
      ge-2/0/8, backup, Physical link is Up
    Bandwidth: 1000mbps
```

show interfaces demux0 (ACI Interface Set Configured)

user@host> **show interfaces demux0.1073741827**

```
  Logical interface demux0.1073741827 (Index 346) (SNMP ifIndex 527)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1802 0x8100.302 ]  Encapsulation:
 ENET2
    Demux: Source Family Inet
    ACI VLAN:
      Dynamic Profile: aci-vlan-set-profile
    Demux:
      Underlying interface: ge-1/0/0 (Index 138)
    Input packets : 18
    Output packets: 16
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re, Unnumbered
      Donor interface: lo0.0 (Index 322)
      Preferred source address: 203.0.113.202
      Addresses, Flags: Primary Is-Default Is-Primary
        Local: 203.0.113.119
    Protocol pppoe
      Dynamic Profile: aci-vlan-pppoe-profile,
      Service Name Table: None,
      Max Sessions: 32000, Max Sessions VSA Ignore: Off,
      Duplicate Protection: On, Short Cycle Protection: Off,
      Direct Connect: Off,
      AC Name: nbc
```

Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

show interfaces extensive

Syntax show interfaces extensive

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers.
Command introduced in Junos OS Release 17.2 for PT1000 and PTX10008 Packet Transport Routers.

Description Display extensive information about all interfaces configured on the router.

Copyright © 2018, Juniper Networks, Inc.

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 is 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.
- The **Output packets** field under the **Traffic statistics** section in the output of the **show interfaces extensive** command includes both IPv4 and IPv6 packets. For example, in a scenario in which both IPv4 and IPv6 packets are being mirrored on the same interface and when you deactivate an IPv4 port-mirroring instance on the chassis, the output of the **show interfaces extensive** command shows a value in the Output packets field of the Traffic statistics section, which is the value of IPv6 packets that are mirrored and not of the IPv4 packets. This behavior is expected.
- For IQ2 PIC interfaces, the output of the **show interfaces extensive** command displays byte statistics that includes Layer 2 headers.
- If there are active OTN defects when an ISSU is performed, and the defect persists after the upgrade completes, the OTN alarm count is incremented by 1. For example, if an OTN alarm is active with a count of 1 and the defect remains after ISSU, the alarm count is incremented to 2. This behavior is expected.

 Copyright © 2018, Juniper Networks, Inc.

Options This command has no options.

Required Privilege Level view

List of Sample Output

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 1110. 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 CE information counters display 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: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  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
```

 Copyright © 2018, Juniper Networks, Inc.

```
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           low    none
  3 network-control       5        5000000   5     0           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
                                           Packets                 Bytes
    Destination class            (packet-per-second)    (bits-per-second)
                      silv1_new                   0                     0
                               (                  0) (                  0)
                      silv2_new                   0                     0
                               (                  0) (                  0)
                      silv_misc                   0                     0
                               (                  0) (                  0)
                        silver0                   0                     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
```

Copyright © 2018, Juniper Networks, Inc.

```
                                (                  0) (                0)
                        silver9                    0                   0
                                (                  0) (                0)
                                             Packets               Bytes
          Source class                 (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: 192.168.220.24/30, Local: 192.168.220.26, Broadcast:
192.168.220.27,         Generation: 150
```

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: 192.168.220.24/30, Local: 192.168.220.26, Broadcast:
192.168.220.27,
        Generation: 170
    Protocol multiservice, MTU: Unlimited, Generation: 156, Route table: 0
      Flags: Is-Primary
      Policer: Input: __default_arp_policer__
```

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
```

 Copyright © 2018, Juniper Networks, Inc.

```
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  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                      Seconds
    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
                            %            bps     %            usec
    0 best-effort          95     9500000000    95               0      low
none
    3 network-control       5      500000000     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                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                  500                    0 pps
     Output packets:                    0                    0 pps
     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: 2001:0db8:0a0b:12f0:0000:0000:0000/112, Local:
2001:0db8:0a0b:12f0:0000:0000:0000:0001
    Generation: 506
      Addresses, Flags: Is-Preferred
        Destination: 0db8::/64, Local: 0db8::21d:b5ff:fef8:6deb
    Protocol multiservice, MTU: Unlimited, Generation: 508
    Generation: 23740, Route table: 0
      Policer: Input: __default_arp_policer__
```

show interfaces extensive (IQ2 and IQ2E)

```
user@host> show interfaces ge-3/2/2 extensive
```

 Copyright © 2018, Juniper Networks, Inc.

```
    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: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
      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 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:                      Receive        Transmit
```

```
      Total octets                      1716096          1716448
      Total packets                       13407            13411
      Unicast packets                     13407            13407
      Broadcast packets                       0                0
      Multicast packets                       0                4
      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
                              %            bps     %           usec
      0 best-effort          95      950000000    95              0      low
  none
      3 network-control       5       50000000     5              0      low
  none
      Direction : Input
      CoS transmit queue               Bandwidth               Buffer Priority
 Limit
                              %            bps     %           usec
      0 best-effort          95      950000000    95              0      low
  none
      3 network-control       5       50000000     5              0      low
  none

    Logical interface ge-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
```

 Copyright © 2018, Juniper Networks, Inc.

```
      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__
```

show interfaces extensive (100-Gigabit Ethernet Type 4 PIC with CFP)

```
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
  Damping      : half-life: 5 sec, max-suppress: 20 sec, reuse 1000, suppress:
2000, state: enabled
  Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  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:                      Receive         Transmit
    Total octets                             0                0
    Total packets                            0                0
    Unicast packets                          0                0
    Broadcast packets                        0                0
    Multicast packets                        0                0
```

 Copyright © 2018, Juniper Networks, Inc.

```
    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       2500000000      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: 192.168.220.24/30, Local: 192.168.220.26, Broadcast:
192.168.220.27, 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
```

show interfaces extensive (PTX5000 Packet Transport Router)

```
user@host> show interfaces et-0/0/6 extensive
Physical interface: et-0/0/6, Enabled, Physical link is Up
  Interface index: 347, SNMP ifIndex: 531, Generation: 350
  Link-level type: Ethernet, MTU: 1514, Speed: 40Gbps, BPDU Error: None, Loop
Detect PDU 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
  Damping         : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0,
```

```
 state: unsuppressed
  Current address: 30:b6:4f:02:29:06, Hardware address: 30:b6:4f:02:29:06
  Last flapped   : 2017-02-15 21:40:06 PST (22:55:13 ago)
  Statistics last cleared: 2017-02-16 20:33:02 PST (00:02:17 ago)
  Traffic statistics:
   Input  bytes  :              1760000                    0 bps
   Output bytes  :              1540000                    0 bps
   Input  packets:                16000                    0 pps
   Output packets:                14000                    0 pps
   IPv6 transit statistics:
    Input  bytes  :              880000
    Output bytes  :              770000
    Input  packets:                8000
    Output packets:                7000
  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                             7000                 7000                    0

    1                                0                    0                    0

    2                                0                    0                    0

    3                             7000                 7000                    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                       2048000          1792000
    Total packets                        16000            14000
    Unicast packets                      16000            14000
    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
```

 Copyright © 2018, Juniper Networks, Inc.

```
    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: 0 (0x00)
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority
Limit
                            %            bps     %            usec
    0 best-effort          95    38000000000    95               0      low
none
    3 network-control       5     2000000000     5               0      low
none
  Preclassifier statistics:
    Traffic Class        Received Packets   Transmitted  Packets      Dropped
Packets
    best-effort                        0                      0
    0
    best-effort                        0                      0
    0
    best-effort                        0                      0
    0
    best-effort                        0                      0
    0
    best-effort                        0                      0
    0
    best-effort                        0                      0
    0
    best-effort                        0                      0
    0
    best-effort                        0                      0
    0
  Link Degrade :
    Link Monitoring                   :  Disable
  Interface transmit statistics: Disabled

  Logical interface et-0/0/6.0 (Index 93) (SNMP ifIndex 841) (Generation 158)
    Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :             1760000
     Output bytes  :             1540000
     Input  packets:               16000
     Output packets:               14000
     IPv6 transit statistics:
      Input  bytes  :             880000
      Output bytes  :             770000
      Input  packets:               8000
      Output packets:               7000
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :             1760000                  0 bps
     Output bytes  :             1540000                  0 bps
     Input  packets:               16000                  0 pps
     Output packets:               14000                  0 pps
```

```
     IPv6 transit statistics:
      Input  bytes  :              880000
      Output bytes  :              770000
      Input  packets:                8000
      Output packets:                7000
    Protocol inet, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new
hold cnt: 0, NH drop cnt: 0
    Generation: 206, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
       Destination: 11.0.0/24, Local: 11.0.0.2, Broadcast: 11.0.0.255, Generation:
 228
    Protocol inet6, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new
hold cnt: 0, NH drop cnt: 0
    Generation: 207, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 1100::/120, Local: 1100::2
    Generation: 230
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::32b6:4fff:fe02:2906
    Protocol multiservice, MTU: Unlimited, Generation: 232
    Generation: 208, Route table: 0
      Policer: Input: __default_arp_policer__
```

sshow interfaces extensive (PTX Routers)

```
user@host> show interfaces ae31 extensive
Physical interface: ae31, Enabled, Physical link is Up
  Interface index: 137, SNMP ifIndex: 511, Generation: 140
  Link-level type: Ethernet, MTU: 1518, Speed: 3Gbps, BPDU Error: None, MAC-REWRITE
 Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled
  Pad to minimum frame size: Disabled
  Minimum links needed: 1, Minimum bandwidth needed: 1bps
  Device flags : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000, LAG_Enhanced
```

show interfaces extensive (PTX10008 Routers)

```
user@host> show interfaces ae0 extensive
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 917, SNMP ifIndex: 817, Generation: 4436
  Link-level type: Ethernet, MTU: 1518, Speed: 20Gbps, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow
control: Disabled, Minimum links needed: 1,
  Minimum bandwidth needed: 1bps
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 30:b6:4f:e9:7c:05, Hardware address: 30:b6:4f:e9:7c:05
  Last flapped   : 2017-04-10 05:20:29 PDT (00:03:52 ago)
  Statistics last cleared: 2017-04-10 05:21:52 PDT (00:02:29 ago)
  Traffic statistics:
   Input  bytes  :          36463816334                    0 bps
   Output bytes  :          36463816334                    0 bps
   Input  packets:             24671053                    0 pps
   Output packets:             24671053                    0 pps
   IPv6 transit statistics:
```


Copyright © 2018, Juniper Networks, Inc.

```
    Input  bytes  :        18231905950
    Output bytes  :        18231905950
    Input  packets:           12335525
    Output packets:           12335525
  MAC statistics:                      Receive        Transmit
    Broadcast packets                        0               0
    Multicast packets                        0               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
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0                         24671053             24671053                    0

    1                                0                    0                    0

    2                                0                    0                    0

    3                                0                    0                    0

  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control

  Logical interface ae0.0 (Index 99) (SNMP ifIndex 832) (Generation 43813)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ]  Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :       4934211          0    7292763858            0
        Output:       4934211          0    7292763858            0
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0
    Link:
      et-0/0/28:0.0
        Input :       4934211          0    7292763858            0
        Output:       4934211          0    7292763858            0
      et-0/0/28:3.0
        Input :             0          0             0            0
        Output:             0          0             0            0


    Aggregate member links: 2

    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      et-0/0/28:0.0              0           0            0            0
      et-0/0/28:3.0              0           0            0            0
    Protocol inet, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new
hold cnt: 0, NH drop cnt: 0
    Generation: 89219, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
       Destination: 21.0.0.0/30, Local: 21.0.0.1, Broadcast: 21.0.0.3, Generation:
```

Copyright © 2018, Juniper Networks, Inc.

```
 62420
    Protocol inet6, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 2, Curr new
hold cnt: 0, NH drop cnt: 0
    Generation: 89220, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 3001::1500:0/126, Local: 3001::1500:1
    Generation: 62422
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::32b6:4f00:2e9:7c05
    Protocol multiservice, MTU: Unlimited, Generation: 62424
    Generation: 89221, Route table: 0
      Policer: Input: __default_arp_policer__

  Logical interface ae0.1 (Index 100) (SNMP ifIndex 833) (Generation 43814)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.3 ]  Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :       4934211          0    7292763858            0
        Output:       4934211          0    7292763858            0
    Adaptive Statistics:
        Adaptive Adjusts:         0
        Adaptive Scans  :         0
        Adaptive Updates:         0
    Link:
      et-0/0/28:0.1
        Input :             0          0             0            0
        Output:       4934211          0    7292763858            0
      et-0/0/28:3.1
        Input :       4934211          0    7292763858            0
        Output:             0          0             0            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      et-0/0/28:0.1            0           0            0            0
      et-0/0/28:3.1            0           0            0            0
    Protocol inet, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new
hold cnt: 0, NH drop cnt: 0
    Generation: 89222, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
       Destination: 21.0.0.4/30, Local: 21.0.0.5, Broadcast: 21.0.0.7, Generation:
 62426
    Protocol inet6, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 2, Curr new
hold cnt: 0, NH drop cnt: 0
    Generation: 89223, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 3001::1500:4/126, Local: 3001::1500:5
    Generation: 62428
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::32b6:4f00:3e9:7c05
    Protocol multiservice, MTU: Unlimited, Generation: 62430
    Generation: 89224, Route table: 0
      Policer: Input: __default_arp_policer__

  Logical interface ae0.2 (Index 101) (SNMP ifIndex 834) (Generation 43815)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.4 ]  Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :       4934211          0    7292763858            0
        Output:       4934211          0    7292763858            0
```

 Copyright © 2018, Juniper Networks, Inc.

```
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0
    Link:
      et-0/0/28:0.2
        Input :       2467106          0    3646382668           0
        Output:       4934211          0    7292763858           0
      et-0/0/28:3.2
        Input :       2467105          0    3646381190           0
        Output:             0          0             0           0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      et-0/0/28:0.2             0           0            0            0
      et-0/0/28:3.2             0           0            0            0
    Protocol inet, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new
hold cnt: 0, NH drop cnt: 0
    Generation: 89225, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 21.0.0.8/30, Local: 21.0.0.9, Broadcast: 21.0.0.11,
Generation: 62432
    Protocol inet6, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new
hold cnt: 0, NH drop cnt: 0
    Generation: 89226, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 3001::1500:8/126, Local: 3001::1500:9
    Generation: 62434
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::32b6:4f00:4e9:7c05
    Protocol multiservice, MTU: Unlimited, Generation: 62436
    Generation: 89227, Route table: 0
      Policer: Input: __default_arp_policer__

  Logical interface ae0.3 (Index 102) (SNMP ifIndex 835) (Generation 43816)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.5 ]  Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :       4934210          0    7292762380           0
        Output:       4934210          0    7292762380           0
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0
    Link:
      et-0/0/28:0.3
        Input :       4934210          0    7292762380           0
        Output:             0          0             0           0
      et-0/0/28:3.3
        Input :             0          0             0           0
        Output:       4934210          0    7292762380           0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      et-0/0/28:0.3             0           0            0            0
      et-0/0/28:3.3             0           0            0            0
    Protocol inet, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new
hold cnt: 0, NH drop cnt: 0
    Generation: 89228, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
```

```
        Destination: 21.0.0.12/30, Local: 21.0.0.13, Broadcast: 21.0.0.15,
Generation: 62438
    Protocol inet6, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 2, Curr new
hold cnt: 0, NH drop cnt: 0
    Generation: 89229, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 3001::1500:c/126, Local: 3001::1500:d
    Generation: 62440
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::32b6:4f00:5e9:7c05
    Protocol multiservice, MTU: Unlimited, Generation: 62442
    Generation: 89230, Route table: 0
      Policer: Input: __default_arp_policer__

  Logical interface ae0.4 (Index 103) (SNMP ifIndex 836) (Generation 43817)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.6 ]  Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :       4934210          0    7292762380            0
        Output:       4934210          0    7292762380            0
    Adaptive Statistics:
        Adaptive Adjusts:         0
        Adaptive Scans  :         0
        Adaptive Updates:         0
    Link:
      et-0/0/28:0.4
        Input :       2467105          0    3646381190            0
        Output:       2467105          0    3646381190            0
      et-0/0/28:3.4
        Input :       2467105          0    3646381190            0
        Output:       2467105          0    3646381190            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      et-0/0/28:0.4            0           0            0            0
      et-0/0/28:3.4            0           0            0            0
    Protocol inet, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new
hold cnt: 0, NH drop cnt: 0
    Generation: 89231, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 21.0.0.16/30, Local: 21.0.0.17, Broadcast: 21.0.0.19,
Generation: 62444
    Protocol inet6, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 2, Curr new
hold cnt: 0, NH drop cnt: 0
    Generation: 89232, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 3001::1500:10/126, Local: 3001::1500:11
    Generation: 62446
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::32b6:4f00:6e9:7c05
    Protocol multiservice, MTU: Unlimited, Generation: 62448
    Generation: 89233, Route table: 0
      Policer: Input: __default_arp_policer__

  Logical interface ae0.32767 (Index 104) (SNMP ifIndex 5645) (Generation 43818)

    Flags: Up SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
```

 Copyright © 2018, Juniper Networks, Inc.

```
    Input :           0          0             0            0
    Output:           0          0             0            0
  Adaptive Statistics:
    Adaptive Adjusts:          0
    Adaptive Scans  :          0
    Adaptive Updates:          0
  Link:
    et-0/0/28:0.32767
    Input :           0          0             0            0
    Output:           0          0             0            0
    et-0/0/28:3.32767
    Input :           0          0             0            0
    Output:           0          0             0            0
  Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
    et-0/0/28:0.32767          0           0            0            0
    et-0/0/28:3.32767          0           0            0            0
  Protocol multiservice, MTU: Unlimited, Generation: 89234, Route table: 0
    Flags: None
    Policer: Input: __default_arp_policer__
```

show interfaces extensive (PTX1000 Routers)

```
user@host> show interfaces et-0/0/48:1 extensive
Physical interface: et-0/0/48:1, Enabled, Physical link is Up
  Interface index: 306, SNMP ifIndex: 697, Generation: 311
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, BPDU Error:
 None, Loop Detect PDU Error: None, MAC-REWRITE 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
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0,
 state: unsuppressed
  Current address: ec:13:db:62:4a:f6, Hardware address: ec:13:db:62:4a:f6
  Last flapped   : 2017-05-08 11:07:59 PDT (12:08:13 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: 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: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0                                0                    0                    0
```

Copyright © 2018, Juniper Networks, Inc.

```
    1                              0                   0                    0

    2                              0                   0                    0

    3                              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                              3
    Errored blocks                          3
  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: 0 (0x00)
  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
  Link Degrade :
    Link Monitoring                   :  Disable
  Interface transmit statistics: Disabled
```

show interfaces extensive (MX Series Routers)

```
user@host> show interfaces xe-0/0/0 extensive
Physical interface: xe-0/0/0, Enabled, Physical link is Up
  Interface index: 145, SNMP ifIndex: 592, Generation: 148
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, BPDU Error:
```

 Copyright © 2018, Juniper Networks, Inc.

```
None,
 Loopback: None, Source filtering: Disabled, Flow control: Enabled
 Pad to minimum frame size: Enabled
 Device flags   : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 Link flags     : None
 CoS queues     : 8 supported, 8 maximum usable queues
 Hold-times     : Up 0 ms, Down 0 ms
 Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
 Last flapped   : 2013-10-26 03:20:40 test (2w3d 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
  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, 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                                0                    0                    0

   1                                0                    0                    0

   2                                0                    0                    0

   3                                0                    0                    0

 Queue number:         Mapped forwarding classes
   0                   best-effort
   1                   expedited-forwarding
   2                   assured-forwarding
   3                   network-control
 Active alarms  : LINK
 Active defects : LINK
 PCS statistics                      Seconds
   Bit errors                           109
   Errored blocks                       109
 MAC statistics:                      Receive         Transmit
   1                                0                    0                    0

   2                                0                    0                    0
```

```
    3                                0                    0                    0

  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control
  Active alarms  : LINK
  Active defects : LINK
  PCS statistics                      Seconds
    Bit errors                           109
    Errored blocks                       109
  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 errors                             0                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: 0
  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
  Interface transmit statistics: Disabled
```

When an ASIC is wedged, the interfaces are brought down along with the IFD. The reason for the link down is displayed as **ASIC-Error** in the **Device flags**.

```
user@host> show interfaces xe-1/0/0 extensive
Physical interface: xe-1/0/0, Administratively down, Physical link is Down
  Interface index: 147, SNMP ifIndex: 563, Generation: 150
  Link-level type: Ethernet, MTU: 1514, MRU: 0, LAN-PHY mode, Speed: 10Gbps, BPDU
 Error: None, Loop Detect PDU Error: None,
  MAC-REWRITE Error: None, Loopback: None, Source filtering: Disabled, Flow
control: Disabled
  Pad to minimum frame size: Disabled
```

 Copyright © 2018, Juniper Networks, Inc.

```
  Device flags   : Present Running Down ASIC-Error
  Interface flags: Hardware-Down Down Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0,
 state: unsuppressed
  Current address: cc:e1:7f:a8:05:4a, Hardware address: cc:e1:7f:a8:05:4a
  Last flapped   : 2017-06-05 17:20:54 PDT (00:03:51 ago)
  Statistics last cleared: Never
```

show interfaces extensive (MX480 Router with MPC5E and 10-Gigabit Ethernet OTN Interface)

```
user@host> show interfaces xe-0/0/3 extensive
Physical interface: xe-0/0/3, Enabled, Physical link is Up
  Interface index: 200, SNMP ifIndex: 577, Generation: 203
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 10Gbps,
BPDU Error: None, MAC-REWRITE Error: None, Loopback: None, Source filtering:
Disabled, Flow control: Enabled
  Pad to minimum frame size: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped   : 2014-06-26 18:16:50 PDT (04:58:35 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
  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, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
  Output errors:
    Carrier transitions: 5, 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                                0                    0                    0

    1                                0                    0                    0

    2                                0                    0                    0
```

```
    3                                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                      4
  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 errors                          0                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: 0 (0x00)
  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
  Interface transmit statistics: Disabled
```

show interfaces extensive (MX480 Router with MPC5E and 100-Gigabit Ethernet OTN Interface)

```
user@host> show interfaces et-2/1/0 extensive
Physical interface: et-2/1/0, Enabled, Physical link is Up
  Interface index: 215, SNMP ifIndex: 872, Generation: 218
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, Speed: 100Gbps, BPDU Error:
None, Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled
  Pad to minimum frame size: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
```

 Copyright © 2018, Juniper Networks, Inc.

```
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped   : 2014-06-26 18:42:04 PDT (04:36:58 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
  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, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
  Output errors:
    Carrier transitions: 263, 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                                0                    0                    0

    1                                0                    0                    0

    2                                0                    0                    0

    3                                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                       754
  MAC statistics:                      Receive         Transmit
    Total octets                         14960                0
    Total packets                          104                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                            2
```

```
    Fragment frames                         6
    VLAN tagged frames                      0
    Code violations                         0
    Total errors                           98                0
  Filter statistics:
    Input packet count                    104
    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
  OTN alarms         : None
  OTN defects        : None
  OTN FEC Mode       : GFEC
  OTN Rate           :     OTU4 100Gbps
  OTN Line Loopback : None
  OTN Local Loopback: None
  OTN Payload PRBS   : None
  OTN FEC statistics:
    Corrected Errors                            169828399453
    Uncorrected Words                            28939961456
    Corrected Error Ratio (      17963 sec average)  8.46e-05
  OTN FEC alarms:          Seconds        Count  State
    FEC Degrade               1180            3  OK
    FEC Excessive             1160            5  OK
  OTN OC:                  Seconds        Count  State
    LOS                        129            1  OK
    LOF                          2            1  OK
    LOM                          0            0  OK
    Wavelength Lock              0            0  OK
  OTN OTU:
    AIS                          0            0  OK
    BDI                          7            1  OK
    IAE                          0            0  OK
    TTIM                       168           45  OK
    BIAE                         0            0  OK
    TSF                          0            0  OK
    SSF                          0            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 ODU:
    AIS                        130            1  OK
    OCI                          0            0  OK
    LCK                          0            0  OK
 BDI                         7            1  OK
    TTIM                       133            1  OK
    IAE                          0            0  OK
    LTC                          0            0  OK
    CSF                          8            4  OK
    TSF                          0            0  OK
    SSF                          0            0  OK
    PTIM                       130            1  OK
    Received DAPI:
```

 Copyright © 2018, Juniper Networks, Inc.

```
    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: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
    Payload Type: 0x00
  ODU Delay Management :
    Result  : 0x00
  PRBS:
    Result: Test not enabled
  OTN Transmitted Overhead Bytes:
    APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
    Payload Type: 0x00
  Packet Forwarding Engine configuration:
    Destination slot: 0 (0x00)
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority
Limit
                              %            bps     %            usec
    0 best-effort            95    95000000000    95               0       low
none
    3 network-control         5     5000000000     5               0       low
none
  Interface transmit statistics: Disabled
```

show interfaces extensive ((MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC)

```
user@host> show interfaces et-3/0/0 extensive
Physical interface: et-3/0/0, Enabled, Physical link is Up
  Interface index: 163, SNMP ifIndex: 564, Generation: 166
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, Speed: 100Gbps, BPDU Error:
None, Loopback: Disabled, Source filtering:
Disabled,
  Flow control: Enabled
  Pad to minimum frame size: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  Wavelength     : 1550.12 nm, Frequency: 193.40 THz
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0,
 state: unsuppressed
  Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped   : 2016-02-17 14:26:31 PST (09:04:28 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
   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, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0,
    L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 5, 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                              0                    0                    0

    1                              0                    0                    0

    2                              0                    0                    0

    3                              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                            8
    Errored blocks                       10
  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 errors                             0                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
```

 Copyright © 2018, Juniper Networks, Inc.

```
  OTN alarms        : None
  OTN defects       : None
  OTN FEC Mode      : SDFEC
  OTN Rate          :     OTU4 (120.5Gbps)
  OTN Line Loopback : None
  OTN Local Loopback: None
  OTN Payload PRBS   : None
  OTN Laser Enable   : On
  OTN FEC statistics:
    Corrected Errors                               7065332638
    Uncorrected Words                                 3412572
    Corrected Error Ratio (      32785 sec average)  1.79e-06 (INVALID)
 OTN FEC alarms:          Seconds        Count  State
    FEC Degrade                 0            0  OK
    FEC Excessive               3            1  OK
  OTN OC:                   Seconds        Count  State
    LOS                         3            1  OK
    LOF                        50            3  OK
    LOM                         3            3  OK
    Wavelength Lock             0            0  OK
  OTN OTU:
    AIS                         0            0  OK
    BDI                         4            4  OK
    IAE                         4            4  OK
    TTIM                        1            1  OK
    BIAE                        3            3  OK
    TSF                        50            3  OK
    SSF                        50            3  OK
    Received DAPI:
    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   ................
    Transmitted DAPI:
    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   ................
  OTN ODU:
    AIS                        20            2  OK
    OCI                         4            4  OK
    LCK                         4            4  OK
    BDI                         2            2  OK
    TTIM                       20            2  OK
    IAE                         0            0  OK
    LTC                         0            0  OK
    CSF                        18            2  OK
    TSF                        66            2  OK
    SSF                        66            2  OK
    PTIM                       43            2  OK
    Received DAPI:
    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   ................
    Transmitted DAPI:
    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   ................
  OTN Received Overhead Bytes:
APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
    Payload Type: 0x07
  ODU Delay Management :
    Result  : 0ms
```

```
  PRBS:
    Result: Test not enabled
  OTN Transmitted Overhead Bytes:
    APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
    Payload Type: 0x07
  Packet Forwarding Engine configuration:
    Destination slot: 0 (0x00)
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority
Limit
                            %            bps     %            usec
    0 best-effort          95    95000000000    95               0       low
none
    3 network-control       5     5000000000     5               0       low
none
  Interface transmit statistics: Disabled
```

show interfaces extensive (PTX3000 Router with 5-port 100-Gigabit DWDM OTN PIC)

```
user@host > show interfaces extensive et-4/0/0
Physical interface: et-4/0/0, Enabled, Physical link is Up
  Interface index: 148, SNMP ifIndex: 544, Generation: 161
  Link-level type: Ethernet, MTU: 1514, Speed: 100Gbps, BPDU Error: None, Loopback:
 Disabled, Source filtering: Disabled,
  Flow control: Enabled
  Device flags   : Present Running
  Interface flags: 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
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0,
 state: unsuppressed
  Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped   : 2016-06-04 21:42:42 PDT (1d 05:09 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: 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: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0                                0                    0                    0

    1                                0                    0                    0
```


Copyright © 2018, Juniper Networks, Inc.

```
  2                             0                    0                    0

  3                             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                             7
Errored blocks                       10
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
OTN alarms       : None
OTN defects      : None
OTN FEC Mode     : SDFEC
OTN Rate         :     OTU4 (120.5Gbps)
OTN Line Loopback : None
OTN Local Loopback: None
OTN Payload PRBS   : None
OTN Laser Enable   : On
OTN FEC statistics:
  Corrected Errors                                   19637746
  Uncorrected Words                                         0
  Corrected Error Ratio (     104923 sec average)  1.55e-09
OTN FEC alarms:          Seconds        Count  State
  FEC Degrade                  0            0  OK
  FEC Excessive                0            0  OK
OTN OC:                  Seconds        Count  State
  LOS                          0            0  OK
  LOF                          2            1  OK
  LOM                          2            1  OK
  Wavelength Lock              0            0  OK
OTN OTU:
  AIS                          0            0  OK
```

```
      BDI                          2            1  OK
      IAE                          0            0  OK
      TTIM                         0            0  OK
      BIAE                         0            0  OK
      TSF                          2            1  OK
      SSF                          0            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 ODU:
      AIS                          0            0  OK
      OCI                          0            0  OK
      LCK                          0            0  OK
      BDI                          2            1  OK
      TTIM                         0            0  OK
      IAE                          0            0  OK
      LTC                          0            0  OK
      CSF                          0            0  OK
      TSF                          2            1  OK
      SSF                          0            0  OK
      PTIM                         2            1  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: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
      Payload Type: 0x07
    ODU Delay Management :
      Result  : 0ms
    PRBS:
      Result: Test not enabled
    OTN Transmitted Overhead Bytes:
      APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
   Payload Type: 0x07
    Packet Forwarding Engine configuration:
      Destination slot: 0 (0x00)
    CoS information:
      Direction : Output
      CoS transmit queue               Bandwidth              Buffer Priority
  Limit
                            %             bps     %            usec
      0 best-effort        95     95000000000    95               0       low
  none
      3 network-control     5      5000000000     5               0       low
  none
    Interface transmit statistics: Disabled
```

show interfaces extensive (MX2020 Router with MPC6E and OTN MIC)

`user@host>` **show interfaces xe-3/0/0 extensive**

 Copyright © 2018, Juniper Networks, Inc.

```
Physical interface: xe-3/0/0, Enabled, Physical link is Up
  Interface index: 166, SNMP ifIndex: 516, Generation: 169
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 10Gbps,
BPDU Error:
None, MAC-REWRITE Error: None, Loopback: None, Source filtering: Disabled, Flow
control:
Enabled
  Pad to minimum frame size: Disabled
  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: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped   : 2014-05-28 17:53:12 PDT (05:56:24 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
  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, 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: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                    0                    0                    0

    1 expedited-forwarding                   0                    0
       0
    2 assured-forwarding                   0                    0
     0
    3 network-control                  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                             2
    Errored blocks                         2
```

```
 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 errors                             0                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
 OTN alarms        : None
 OTN defects       : None
 OTN FEC Mode      : GFEC
 OTN Rate          : Fixed Stuff Bytes 11.0957Gbps
 OTN Line Loopback : None
 OTN Local Loopback: None
 OTN Payload PRBS   : None
 OTN FEC statistics:
   Corrected Errors                                      0
   Uncorrected Words                                     0
   Corrected Error Ratio (      21387 sec average)  0.00e+00
 OTN FEC alarms:         Seconds        Count  State
   FEC Degrade                 0            0  OK
   FEC Excessive               0            0  OK
 OTN OC:                 Seconds        Count  State
   LOS                         0            0  OK
   LOF                         0            0  OK
   LOM                         0            0  OK
   Wavelength Lock             0            0  OK
 OTN OTU:
   AIS                         0            0  OK
   BDI                         0            0  OK
   IAE                         0            0  OK
   TTIM                        0            0  OK
   BIAE                        0            0  OK
   TSF                         0            0  OK
   SSF                         0            0  OK
   Received DAPI:
   00 53 4d 2d 54 52 43 20 44 41 50 49 2d 53 45 43   .SM-TRC DAPI-SEC
   Received SAPI:
   00 53 4d 2d 54 52 43 20 53 41 50 49 2d 53 45 43   .SM-TRC SAPI-SEC
   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 ODU:
```

 Copyright © 2018, Juniper Networks, Inc.

```
    AIS                         0            0  OK
    OCI                         0            0  OK
 LCK                         0            0  OK
    BDI                         0            0  OK
    TTIM                        0            0  OK
    IAE                         0            0  OK
    LTC                         0            0  OK
    CSF                         0            0  OK
    TSF                         0            0  OK
    SSF                         0            0  OK
    PTIM                        0            0  OK
    Received DAPI:
    00 50 4d 2d 54 52 43 20 44 41 50 49 2d 53 45 43   .PM-TRC DAPI-SEC
    Received SAPI:
    00 50 4d 2d 54 52 43 20 53 41 50 49 2d 53 45 43   .PM-TRC SAPI-SEC
    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: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
    Payload Type: 0x00
  ODU Delay Management :
    Result  : 0x00
  PRBS:
    Result: Test not enabled
  OTN Transmitted Overhead Bytes:
    APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
    Payload Type: 0x00
  Packet Forwarding Engine configuration:
    Destination slot: 0 (0x00)
  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
  Interface transmit statistics: Disabled
```

show interfaces extensive (MX2010 Router with MPC6E and 100-Gigabit Ethernet OTN Interface)

user@host> **show interfaces et-9/0/0 extensive**

```
Physical interface: et-9/0/0, Enabled, Physical link is Up
  Interface index: 196, SNMP ifIndex: 623, Generation: 199
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, Speed: 100Gbps, BPDU Error:
None, Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled
  Pad to minimum frame size: Disabled
  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: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped   : 2014-06-26 18:18:34 PDT (04:17:07 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
    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, 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                             0                    0                    0

      1                             0                    0                    0

      2                             0                    0                    0

      3                             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
      Total errors                             0                0
    Filter statistics:
      Input packet count                       0
      Input packet rejects                     0
      Input DA rejects                         0
```

 Copyright © 2018, Juniper Networks, Inc.

```
    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: 0 (0x00)
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority
Limit
                              %            bps     %            usec
    0 best-effort            95    95000000000    95               0      low
none
    3 network-control         5     5000000000     5               0      low
none
  Interface transmit statistics: Disabled
```

show interfaces extensive (MX2010 Router with MPC6E and 10-Gigabit Ethernet Interface)

```
user@host> show interfaces xe-6/1/0 extensive
Physical interface: xe-6/1/0, Enabled, Physical link is Up
  Interface index: 159, SNMP ifIndex: 603, Generation: 162
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 10Gbps,
BPDU Error: None, MAC-REWRITE Error: None, Loopback: None, Source filtering:
Disabled, Flow control: Enabled
  Pad to minimum frame size: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped   : 2014-06-26 18:16:50 PDT (04:21:04 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
  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, 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                                0                    0                      0

     1                                0                    0                      0

     2                                0                    0                      0

     3                                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                         1
   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 errors                             0                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: 0 (0x00)
   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
   Interface transmit statistics: Disabled
```

 Copyright © 2018, Juniper Networks, Inc.

show interfaces extensive (T4000 Routers with Type 5 FPCs)

The output fields for the **show interfaces *interface* extensive** command remains the same for 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP), 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-24XGE-SFPP), and 100-Gigabit Ethernet Type 5 PIC with CFP (PF-1CGE-CFP).

```
user@host> show interfaces xe-4/0/0 extensive
Physical interface: xe-4/0/0, Enabled, Physical link is Up
  Interface index: 200, SNMP ifIndex: 592, Generation: 203
  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
  Damping      : half-life: 5 sec, max-suppress: 20 sec, reuse 1000, suppress:
2000, state: enabled
  Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped   : 2013-06-03 16:01:56 PDT (06:04:07 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: 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 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
```

Copyright © 2018, Juniper Networks, Inc.

```
  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: 0 (0x00)
  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

   real-time                          0                    0                    0
   network-control                    0                    0                    0
   best-effort                        0                    0                    0
  Interface transmit statistics: Disabled
```

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: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  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
```


Copyright © 2018, Juniper Networks, Inc.

```
    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:        Role              System              System         Port
     Port  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
      Policer: Input: __default_arp_policer__
```

 Copyright © 2018, Juniper Networks, Inc.

show interfaces lsi (Label-Switched Interface)

Syntax

show interfaces *interface-type*
<brief | detail | extensive | terse>
<descriptions>
<media>
<routing-instance *instance-name*>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display status information about the specified label-switched interface (LSI).

Options

interface-type—On most routers, the interface type is **lt-*fpc/pic/port***.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

routing-instance *instance-name*—(Optional) Display information for the specified routing instance.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Related Documentation

List of Sample Output

show interfaces lsi extensive on page 1601

Output Fields

Table 78 on page 1599 lists the output fields for the **show interfaces** (logical tunnel) command. Output fields are listed in the approximate order in which they appear.

Table 78: Logical Tunnel show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Logical Interface		

Table 78: Logical Tunnel show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	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 1110.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. • **Input bytes**—Rate of bytes received on the interface. • **Output bytes**—Rate of bytes transmitted on the interface. • **Input packets**—Rate of packets received on the interface. • **Output packets**—Rate of packets transmitted on the interface.	**detail extensive**
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	**detail extensive**
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	**detail extensive**
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**
Flags	Information about the protocol family flags. Possible values are described in the “Family Flags” section under “Common Output Fields Description” on page 1110.	**detail extensive** none

 Copyright © 2018, Juniper Networks, Inc.

Sample Output

show interfaces lsi extensive

```
user@host> show interfaces lsi extensive
Physical interface: lsi

  Logical interface lsi.84934656 (Index 363) (SNMP ifIndex 586) (Generation 194)

    Flags: Up Point-To-Point SNMP-Traps 0x4000000 Encapsulation: LSI-NULL
    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 vpls, MTU: Unlimited, Generation: 279, Route table: 10

  Logical interface lsi.84934657 (Index 366) (SNMP ifIndex 589) (Generation 197)

    Flags: Up Point-To-Point SNMP-Traps 0x4000000 Encapsulation: LSI-NULL
    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 vpls, MTU: Unlimited, Generation: 282, Route table: 10
```

Copyright © 2018, Juniper Networks, Inc.

show interfaces media

Syntax show interfaces media

Release Information Command introduced before Junos OS Release 7.4.
Command introduced on PTX Series Packet Transport Routers for Junos OS Release 12.1.

Description Display media-specific information about all configured network interfaces.

NOTE: **show interfaces media** lists details for all interfaces, whereas **show interfaces media** ***interface-name*** lists details only for the specified interface.

Options This command has no options.

Additional Information Output from both the **show interfaces** ***interface-name*** **detail** and the **show interfaces** ***interface-name*** **extensive** commands includes all the information displayed in the output from the **show interfaces media** command.

Required Privilege Level view

List of Sample Output
show interfaces media (SONET/SDH) on page 1603
show interfaces media (MX Series Routers) on page 1603
show interfaces media (PTX Series Packet Transport Routers) on page 1604

Output Fields The output from the **show interfaces media** command includes fields that display interface media-specific information. These fields are also included in the **show interfaces** ***interface-name*** command for each particular interface type, and the information provided in the fields is unique to each interface type.

One field unique to the **show interfaces media** command is **interface-type** errors (for example, **SONET errors**). This field appears for channelized E3, channelized T3, channelized OC, E1, E3, SONET, T1, and T3 interfaces. The information provided in this output field is also provided in the output from the **show interfaces** ***interface-name*** command. (For example, for SONET interfaces, these fields are **SONET section**, **SONET line**, and **SONET path**). For a description of errors, see the chapter with the particular interface type in which you are interested.

 Copyright © 2018, Juniper Networks, Inc.

Sample Output

show interfaces media (SONET/SDH)

The following example displays the output fields unique to the **show interfaces media** command for a SONET interface (with no level of output specified):

```
user@host> show interfaces media so-4/1/2
Physical interface: so-4/1/2, Enabled, Physical link is Up
  Interface index: 168, SNMP ifIndex: 495
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 16384
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 1783 (00:00:00 ago), Output: 1786 (00:00:08 ago)
  LCP state: Opened
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
  mpls: Not-configured
  CHAP state: Not-configured
  CoS queues     : 8 supported
  Last flapped   : 2005-06-15 12:14:59 PDT (04:31:29 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : None
  SONET defects  : None
  SONET errors:
     BIP-B1: 121, BIP-B2: 916, REI-L: 0, BIP-B3: 137, REI-P: 16747, BIP-BIP2: 0
  Received path trace: routerb so-1/1/2
  Transmitted path trace: routera so-4/1/2
```

show interfaces media (MX Series Routers)

```
user@host>show interfaces media xe-0/0/0
Physical interface: xe-0/0/0, Enabled, Physical link is Up
  Interface index: 145, SNMP ifIndex: 592
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, BPDU Error:
 None,
  Loopback: None, Source filtering: Disabled, Flow control: Enabled
  Pad to minimum frame size: Enabled
  Device flags   : Present Running
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 08:81:f4:82:a3:f0, Hardware address: 08:81:f4:82:a3:f0
  Last flapped   : 2013-10-26 03:20:40 test (1w6d 00:19 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK
  PCS statistics                     Seconds
    Bit errors                           78
    Errored blocks                       78
  MAC statistics:
    Input bytes: 0, Input packets: 0, Output bytes: 0, Output packets: 0
  Filter statistics:
    Filtered packets: 0, Padded packets: 0, Output packet errors: 0
  Interface transmit statistics: Disabled
```

Copyright © 2018, Juniper Networks, Inc.

show interfaces media (PTX Series Packet Transport Routers)

```
user@host> show interfaces media em0
Physical interface: em0, Enabled, Physical link is Up
      Interface index: 8, SNMP ifIndex: 0
      Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps
      Device flags   : Present Running
      Interface flags: SNMP-Traps
      Link type      : Full-Duplex
      Current address: 00:80:f9:25:00:1b, Hardware address: 00:80:f9:25:00:1b
      Last flapped   : Never
    Input packets : 215151
    Output packets: 72
```


Copyright © 2018, Juniper Networks, Inc.

show interfaces terse

Syntax show interfaces terse

Release Information Command introduced before Junos OS Release 7.4.
Command introduced on PTX Series Packet Transport Routers for Junos OS Release 12.1.

Description Display summary information about interfaces.

Options This command has no options.

Additional Information Interfaces are always displayed in numerical order, from the lowest to the highest FPC slot number. Within that slot, the lowest PIC slot is shown first. On an individual PIC, the lowest port number is always first.

Required Privilege Level view

Related Documentation
- *Examples: Using Logical Systems*

List of Sample Output
show interfaces terse on page 1606
show interfaces terse (TX Matrix Plus Router) on page 1606
show interfaces terse (PTX Series Packet Transport Routers) on page 1607

Output Fields Table 79 on page 1605 lists the output fields for the **show interfaces terse** command. Output fields are listed in the approximate order in which they appear.

Table 79: show interfaces terse Output Fields

Field Name	Field Description
Interface	Interface name.
Admin	Whether the interface is turned on (up) or off (down).
Link	Link state: **up** or **down**.
Proto	Protocol family configured on the logical interface. A logical interface on a router that supports Ethernet OAM always shows the multiservice protocol.
Local	Local IP address of the logical interface.
Remote	Remote IP address of the logical interface.

Sample Output

show interfaces terse

```
user@host> show interfaces terse
Interface               Admin Link Proto    Local                 Remote
t1-0/1/0:0              up    up
t1-0/1/0:0.0            up    up   inet     192.168.220.18/30
t1-0/1/0:1              up    up
t1-0/1/0:2              up    up
t1-0/1/0:3              up    up
at-1/0/0                up    up
at-1/0/1                up    up
dsc                     up    up
fxp0                    up    up
fxp0.0                  up    up   inet     192.168.71.249/21
fxp1                    up    up
fxp1.0                  up    up   inet     10.0.0.4/8
                                   tnp      4
gre                     up    up
ipip                    up    up
lo0                     up    up
lo0.0                   up    up   inet     10.0.1.4            --> 0/0
                                            127.0.0.1           --> 0/0
lo0.16385               up    up   inet

lsi                     up    up
mtun                    up    up
```

show interfaces terse (TX Matrix Plus Router)

```
user@host> show interfaces terse

Interface               Admin Link Proto    Local                 Remote
xe-0/0/0                up    up
xe-0/0/1                up    up
xe-0/0/2                up    up
xe-0/0/3                up    up
xe-6/0/0                up    up
xe-6/0/1                up    up
xe-6/0/2                up    up
xe-6/0/3                up    up
xe-6/1/0                up    up
xe-6/1/1                up    up
xe-6/1/2                up    up
xe-6/1/3                up    up
so-0/0/0                up    up
so-0/0/0.0              up    up   inet     1.1.1.1/30
ge-1/3/0.0              up    up   inet     --> 0/0
ge-7/0/0                up    up
ge-7/0/0.0              up    up   inet     2.15.1.1/30
ge-7/0/0.1              up    up   inet     2.15.1.5/30
ge-7/0/0.2              up    up   inet     2.15.1.9/30
ge-7/0/0.3              up    up   inet     2.15.1.13/30
ge-7/0/0.4              up    up   inet     2.15.1.17/30
ge-7/0/0.5              up    up   inet     2.15.1.21/30
...
em0                     up    up
em0.0                   up    up   inet     192.168.178.11/25
```


Copyright © 2018, Juniper Networks, Inc.

```
gre                     up    up
ipip                    up    up
ixgbe0                  up    up
ixgbe0.0                up    up   inet      10.34.0.4/8
                                             162.0.0.4/2
                                   inet6     fe80::200:ff:fe22:4/64
                                             fec0::a:22:0:4/64
                                   tnp       0x22000004
ixgbe1                  up    up
ixgbe1.0                up    up   inet      10.34.0.4/8
                                             162.0.0.4/2
                                   inet6     fe80::200:1ff:fe22:4/64
                                             fec0::a:22:0:4/64
                                   tnp       0x22000004
```

show interfaces terse (PTX Series Packet Transport Routers)

```
user@host> show interfaces em0 terse
   Interface                Admin Link Proto    Local                 Remote
    em0                      up    up
    em0.0                    up    up   inet     192.168.3.30/24
```

Copyright © 2018, Juniper Networks, Inc.

CHAPTER 15

ANCP Operational Commands

clear ancp neighbor

Syntax

clear ancp neighbor
<ip-address *ip-address*>
<system-name *mac-address*>

Release Information

Command introduced in Junos OS Release 9.4.

Description

Clear the ANCP agent connection with all ANCP neighbors or with the specified ANCP neighbor. This command deletes information for subscribers associated with the neighbor, causing the adjusted traffic rates to revert to the configured rate for the subscriber interfaces. The neighbor remains configured (its administrative state is *enabled*) and can reestablish adjacencies.

This command initiates logout of ANCP-triggered dynamic VLAN sessions on the physical interface associated with the specified neighbor; conventionally autosensed dynamic VLAN sessions and their associated logical interfaces are not affected.

Options

none—Clear all ANCP neighbors.

ip-address *ip-address*—(Optional) Clear the ANCP neighbor specified by the IP address.

system-name *mac-address*—(Optional) Clear the ANCP neighbor specified by the MAC address.

Required Privilege Level

clear

Related Documentation

- show ancp neighbor on page 1624

List of Sample Output

clear ancp neighbor on page 1610
show ancp neighbor on page 1610

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the **show ancp neighbor** command before and after clearing the ANCP neighbors to verify the clear operation.

Sample Output

clear ancp neighbor

```
user@host> clear ancp neighbor
```

show ancp neighbor

The following sample output displays the connections with ANCP neighbors before and after the **clear ancp neighbor** command was issued.

Copyright © 2018, Juniper Networks, Inc.

```
user@host> show ancp neighbor

  IP Address         MAC Address         State          Subscriber   Capabilities
                                                          Count
  203.0.113.102      00:00:5e:00:53:10  Established    5              Topo
  203.0.113.122      00:00:5e:00:53:12  Established    5              Topo
  203.0.113.132      00:00:5e:00:53:13  Established    5              Topo
  203.0.113.142      00:00:5e:00:53:14  Established    5              Topo
```

```
user@host> clear ancp neighbor ip-address 203.0.113.102
```

```
user@host> show ancp neighbor

  IP Address         MAC Address         State          Subscriber   Capabilities
                                                          Count
  203.0.113.122      00:00:5e:00:53:12  Established    5              Topo
  203.0.113.132      00:00:5e:00:53:13  Established    5              Topo
  203.0.113.142      00:00:5e:00:53:14  Established    5              Topo
```

clear ancp statistics

Syntax

clear ancp statistics
<ip-address *ip-address*>
<system-name *mac-address*>

Release Information Command introduced in Junos OS Release 13.3.

Description Clear current statistics accumulated by the ANCP agent for all ANCP neighbors or the specified neighbor.

Options **none**—Clear all ANCP statistics.

ip-address *ip-address*—(Optional) Clear statistics for the ANCP neighbor specified by the IP address.

system-name *mac-address*—(Optional) Clear statistics for the ANCP neighbor specified by the MAC address.

Required Privilege Level clear

Related Documentation

- show ancp neighbor on page 1624

List of Sample Output

clear ancp statistics on page 1612
show ancp neighbor on page 1612

Output Fields When you enter this command, you are provided no feedback on the status of your request. You can enter the **show ancp neighbor** command before and after clearing the ANCP neighbor statistics to verify the clear operation.

Sample Output

clear ancp statistics

```
user@host> clear ancp statistics
```

show ancp neighbor

The following sample output displays statistics for an ANCP neighbor before and after the **clear ancp statistics** command was issued.

```
user@host> show ancp neighbor ip-address 192.168.10.1 detail
Neighbor Information
      IP Address              : 192.168.10.1
      System Name             : 00:00:5E:00:53:02
          Up Time                             : 38
          TCP Port                            : 64959
```

 Copyright © 2018, Juniper Networks, Inc.

```
          State                           : Established
          Subscriber Count                : 7
          Capabilities                    : Topology Discovery
          System Instance                 : 11
          Peer Instance                   : 1
          Adjacency Timer (in 100ms)      : 50
          Peer Adjacency Timer (in 100ms) : 100
          Partition Type                  : 0
          Partition Flag                  : 1
          Partition Identifier            : 0
          Dead Timer                      : 22
          Received Syn Count              : 47
          Received Synack Count           : 48
          Received Rstack Count           : 2
          Received Ack Count              : 12
          Received Port Up Count          : 8
          Received Port Down Count        : 2
          Received Other Count            : 0
          Sent Syn Count                  : 48
          Sent Synack Count               : 47
          Sent Rstack Count               : 1
          Sent Ack Count                  : 12
          Max Discovery Limit Exceed Count    : 0
```

```
user@host> clear ancp statistics ip-address 192.168.10.1
```

```
user@host> show ancp neighbor ip-address 192.168.10.1 detail
Neighbor Information
      IP Address              : 192.168.10.1
      System Name             : 00:00:5E:00:53:02
          Up Time                         : 38
          TCP Port                        : 64959
          State                           : Established
          Subscriber Count                : 7
          Capabilities                    : Topology Discovery
          System Instance                 : 11
          Peer Instance                   : 1
          Adjacency Timer (in 100ms)      : 50
          Peer Adjacency Timer (in 100ms) : 100
          Partition Type                  : 0
          Partition Flag                  : 1
          Partition Identifier            : 0
          Dead Timer                      : 22
          Received Syn Count              : 0
          Received Synack Count           : 0
          Received Rstack Count           : 0
          Received Ack Count              : 0
          Received Port Up Count          : 0
          Received Port Down Count        : 0
          Received Other Count            : 0
          Sent Syn Count                  : 0
          Sent Synack Count               : 0
          Sent Rstack Count               : 0
          Sent Ack Count                  : 0
          Max Discovery Limit Exceed Count    : 0
```

clear ancp subscriber

Syntax

clear ancp subscriber
<identifier *identifier*>
<ip-address *ip-address*>
<system-name *mac-address*>

Release Information Command introduced in Junos OS Release 11.4.

Description Clear the ANCP agent connection with all ANCP subscribers or with the specified ANCP subscriber. This command deletes information for the subscribers, causing the adjusted traffic rate to revert to the configured rate for the subscriber interface, but otherwise has no affect on ANCP neighbors.

Options **none**—Clear all ANCP subscribers.

identifier *identifier-string*—(Optional) Clear the ANCP subscriber identified by the access loop ID.

ip-address *ip-address*—(Optional) Clear all ANCP subscribers on the neighbor specified by the IP address.

system-name *mac-address*—(Optional) Clear all ANCP subscribers on the neighbor specified by the MAC address.

Required Privilege Level clear

Related Documentation

- show ancp subscriber on page 1637

List of Sample Output show ancp subscriber brief on page 1614
clear ancp subscriber on page 1615

Output Fields When you enter this command, you are provided no feedback on the status of your request. You can enter the **show ancp subscriber** command before and after clearing the ANCP neighbors to verify the clear operation.

Sample Output

show ancp subscriber brief

```
user@host> show ancp subscriber brief
Loop Identifier          Type      Interface        Rate          Neighbor
                                                     Kbps
  port-1-10               VDSL2     set-ge-10410      64        203.0.113.102

  port-1-11               VDSL2     set-ge-10411      64        203.0.113.112

  port-2-10               VDSL2     ge-1/0/4.12       64        203.0.113.122
```

 Copyright © 2018, Juniper Networks, Inc.

```
  port-2-10                VDSL2    ge-1/0/4.12        64       203.0.113.123
  port-2-11                VDSL2    ge-1/0/4.13        64       203.0.113.132
```

`user@host>` **clear ancp subscriber identifier port-2-10**

`user@host>` **show ancp subscriber brief**

```
Loop Identifier          Type      Interface        Rate          Neighbor
                                                      Kbps
  port-1-10                VDSL2    set-ge-10410       64       203.0.113.102

  port-1-11                VDSL2    set-ge-10411       64       203.0.113.112

  port-2-11                VDSL2    ge-1/0/4.13        64       203.0.113.132
```

clear ancp subscriber

`user@host>` **clear ancp subscriber**

request ancp oam interface

Syntax

request ancp oam interface
(*interface-name* | interface-set *set-name*)
<count *count*>
<timeout *duration*>

Release Information Command introduced in Junos OS Release 11.4.

Description Trigger the access node to run a loopback test on the local loop between the access node and the customer premises equipment. You must specify either an ANCP interface or an ANCP interface set. The access node responds to the NAS with the results of the test.

Options

interface-name—Name of the ANCP interface on whose local loop the loopback test is run.

interface-set *set-name*—Name of the ANCP interface set on whose local loop the loopback test is run.

count *count*—(Optional) Number of times a loopback message is sent on the local loop.
Range: **1** through **32**. Default: **1**.

timeout *duration*—(Optional) Period of time in seconds that the NAS waits for a response to the OAM request.
Range**: 0** through **255**. Default: **5**.

Required Privilege Level view

Related Documentation

- *Triggering ANCP OAM to Test the Local Loop*

List of Sample Output request ancp oam interface on page 1616

Output Fields When you enter this command, you are provided feedback on the status of your request, including the result of the test, the response code, and the response string returned with the OAM response in the event of failure, an error code is displayed.

Sample Output

request ancp oam interface

```
user@host> request ancp oam interface ge-1/0/4.12 count 5 timeout 40
request succeeded
 0x503 : DSL line status showtime
 DEFAULT RESPONSE
```


Copyright © 2018, Juniper Networks, Inc.

request ancp oam neighbor

Syntax

request ancp oam neighbor
(ip-address ***ip-address*** **| system-name** ***neighbor-name*****)**
subscriber ***identifier-string***
<count ***count*****>**
<timeout ***duration*****>**

Release Information

Command introduced in Junos OS Release 11.4.

Description

Trigger the access node to run a loopback test on the local loop between the access node and the customer premises equipment. You must specify both the access node and the subscriber. The access node responds to the NAS with the results of the test.

Options

ip-address ***ip-address***—IP address that specifies the access node on whose local loop the loopback test is run.

system-name ***neighbor-name***—System name that specifies the access node on whose local loop the loopback test is run.

subscriber ***identifier-string***—Access identifier that specifies the subscriber on whose local loop the loopback test is run.

count ***count***—(Optional) Number of times a loopback message is sent on the local loop.
Range: **1** through **32**. Default: **1**.

timeout ***duration***—(Optional) Period of time in seconds that the NAS waits for a response to the OAM request.
Range: **0** through **255**. Default: **5**.

Required Privilege Level

view

Related Documentation

- *Triggering ANCP OAM to Test the Local Loop*

List of Sample Output

request ancp oam subscriber on page 1617

Output Fields

When you enter this command, you are provided feedback on the status of your request, including the result of the test, the response code, and the response string returned with the OAM response in the event of failure, an error code is displayed.

Sample Output

request ancp oam subscriber

`user@host>` **request ancp oam neighbor 203.0.113.21 subscriber “dslam port-1-11”**

Copyright © 2018, Juniper Networks, Inc.

```
request succeeded
 0x503 : DSL line status showtime
 DEFAULT RESPONSE
```

 Copyright © 2018, Juniper Networks, Inc.

show ancp cos

Syntax

show ancp cos
<identifier *identifier*>
<last-update>
<pending-update>

Release Information

Command introduced in Junos OS Release 9.4.

Description

Display information about the CoS state for subscriber traffic.

Options

identifier *identifier*—(Optional) Display information about the local loops for the specified access identifier.

last-update—(Optional) Display the most recently updated CoS information.

pending-update—(Optional) Display the pending update of CoS information.

Required Privilege Level

view

Related Documentation

- show ancp neighbor on page 1624
- show ancp statistics on page 1632
- show ancp subscriber on page 1637

List of Sample Output

show ancp cos on page 1621
show ancp cos last-update on page 1622
show ancp cos pending-update on page 1623

Output Fields

Table 80 on page 1619 lists the output fields for the **show ancp cos** command. Output fields are listed in the approximate order in which they appear.

Table 80: show ancp cos Output Fields

Field Name	Field Description
Per-DSL CoS adjustment	Adjustment values applied by the ANCP agent to the actual downstream rates and frame overhead for frame-mode DSL types. The agent then reports the adjusted rates to CoS to establish a shaping rate for the CoS node that corresponds to the subscriber access line.
QoS Adjust Flag	State of QoS adjust: • **TRUE**—The ANCP agent is enabled to adjust the actual downstream data rates and frame overhead and report the adjusted values to CoS. • **FALSE**—The ANCP agent is not enabled to adjust and report values to CoS.

Copyright © 2018, Juniper Networks, Inc.

Table 80: show ancp cos Output Fields (continued)

Field Name	Field Description
ADSL bytes	Number of bytes by which the actual ADSL downstream cell overhead is adjusted before reporting it to CoS.
ADSL2 bytes	Number of bytes by which the actual ADSL2 downstream cell overhead is adjusted before reporting it to CoS.
ADSL2-PLUS bytes	Number of bytes by which the actual ADSL2+ downstream cell overhead is adjusted before reporting it to CoS.
SDSL overhead adjusted	Percentage by which the actual SDSL downstream rate is adjusted before reporting it to CoS.
SDSL bytes	Number of bytes by which the actual SDSL downstream frame overhead is adjusted before reporting it to CoS.
OTHER overhead adjusted	Percentage by which the actual OTHER downstream rate is adjusted before reporting it to CoS.
OTHER bytes	Number of bytes by which the actual OTHER downstream frame overhead is adjusted before reporting it to CoS.
VDSL overhead adjusted	Percentage by which the actual VDSL downstream rate is adjusted before reporting it to CoS.
VDSL bytes	Number of bytes by which the actual VDSL downstream frame overhead is adjusted before reporting it to CoS.
VDSL2 overhead adjusted	Percentage by which the actual VDSL2 downstream rate is adjusted before reporting it to CoS.
VDSL2 bytes	Number of bytes by which the actual VDSL2 downstream frame overhead is adjusted before reporting it to CoS.
Per-DSL adjustment for reporting	Adjustment values applied by the ANCP agent to the actual downstream rates for individual DSL types to account for traffic overhead. The agent then reports the adjusted rates to AAA.
ADSL adjustment factor	Percentage by which the actual ADSL downstream rate is adjusted before reporting it to AAA.
ADSL2 adjustment factor	Percentage by which the actual ADSL2 downstream rate is adjusted before reporting it to AAA.
ADSL2+ adjustment factor	Percentage by which the actual ADSL2+ downstream rate is adjusted before reporting it to AAA.
VDSL adjustment factor	Percentage by which the actual VDSL downstream rate is adjusted before reporting it to AAA.

Copyright © 2018, Juniper Networks, Inc.

Table 80: show ancp cos Output Fields (continued)

Field Name	Field Description
VDSL2 adjustment factor	Percentage by which the actual VDSL2 downstream rate is adjusted before reporting it to AAA.
SDSL adjustment factor	Percentage by which the actual SDSL downstream rate is adjusted before reporting it to AAA.
OTHER adjustment factor	Percentage by which the actual OTHER downstream rate is adjusted before reporting it to AAA.
Keepalive Timer	Interval between the keepalive messages that the ANCP agent sends to CoS.
Cos State	State of the interaction between the ANCP agent and CoS: • ANCPD_COS_CONNECT_NEEDED • ANCPD_COS_CONNECT_PENDING • ANCPD_COS_CONNECT_DONE • ANCPD_COS_SESSION_SENT • ANCPD_COS_WRITE_READY
Connect Time	Time at which the ANCP agent connected to CoS; useful for debugging.
Session Time	Time at which the ANCP agent sent a session connect message to CoS; useful for debugging.
Routing Instance Time	Time at which the ANCP agent sent the routing instance to CoS; useful for debugging.
Keepalive Time	Time at which the last keepalive message was sent.
Update Time	Time at which the shaping rate was last updated.
Type	Subscriber access type: **ifl** indicates that a single VLAN carries subscriber traffic and **iflset** indicates that a set of VLANs carries subscriber traffic.
Name	System-wide name of the particular subscriber access.
Index	Access identifier.
Pending Update	Actual downstream data rate to be applied next to this local loop, in Kbps.
Last Update	Adjusted downstream data rate last reported to CoS by the ANCP agent for this local loop, in Kbps.

Sample Output

show ancp cos

`user@host>` **show ancp cos**

```
Per-DSL CoS adjustment:
  Qos Adjust Flag:         TRUE
  ADSL bytes:              20
  ADSL2 bytes:              20
  ADSL2-PLUS bytes:              20
  VDSL overhead adjusted:  90
  VDSL bytes:              20
  VDSL2 overhead adjusted: 95
  VDSL2 bytes:             -20
  SDSL overhead adjusted:  85
  SDSL bytes:              30
  OTHER overhead adjusted:  85
  OTHER bytes:              30

Per-DSL adjustment for reporting:
  ADSL adjustment factor:  100
  ADSL2 adjustment factor: 100
  ADSL2+ adjustment factor:100
  VDSL adjustment factor:  100
  VDSL2 adjustment factor: 100
  SDSL adjustment factor:  100
  OTHER adjustment factor:  100


Keepalive Timer:         45 secs
State:                   WRITE_READY
Connect Time:            Fri May  2 12:08:49 2016
Session Time:            Fri May  2 12:18:52 2016
Routing Instance Time:   Fri May  2 12:18:53 2016
Keepalive Time:          Fri May  2 13:44:14 2016
Update Time:             Fri May  2 13:02:55 2016


Type      Name              Index         Pending Update      Last Update
iflset    aci-1004-ge-2/0/0.1073741834 4 None                  36000 Kbps
```

show ancp cos last-update

```
Per-DSL CoS adjustment:
  Qos Adjust Flag:         TRUE
  ADSL bytes:              20
  ADSL2 bytes:              20
  ADSL2-PLUS bytes:              20
  VDSL overhead adjusted:  90
  VDSL bytes:              20
  VDSL2 overhead adjusted: 95
  VDSL2 bytes:             -20
  SDSL overhead adjusted:  85
  SDSL bytes:              30
  OTHER overhead adjusted:  85
  OTHER bytes:              30

Per-DSL adjustment for reporting:
  ADSL adjustment factor:  100
  ADSL2 adjustment factor: 100
  ADSL2+ adjustment factor:100
  VDSL adjustment factor:  100
  VDSL2 adjustment factor: 100
  SDSL adjustment factor:  100
  OTHER adjustment factor:  100
```


Copyright © 2018, Juniper Networks, Inc.

```
Keepalive Timer:         45 secs
State:                   WRITE_READY
Connect Time:            Fri May  2 12:08:49 2016
Session Time:            Fri May  2 12:18:52 2016
Routing Instance Time:   Fri May  2 12:18:53 2016
Keepalive Time:          Fri May  2 13:44:34 2016
Update Time:             Fri May  2 13:02:55 2016


Type      Name             Index         Pending Update      Last Update
iflset    aci-1004-ge-2/0/0.1073741834 4 None                 36000 Kbps
```

show ancp cos pending-update

```
user@host> show ancp cos pending-update
Per-DSL CoS adjustment:
  Qos Adjust Flag:         TRUE
  VDSL overhead adjusted:  90
  VDSL bytes:              20
  VDSL2 overhead adjusted: 95
  VDSL2 bytes:             -20
  SDSL overhead adjusted:  85
  SDSL bytes:              30
  OTHER overhead adjusted:  85
  OTHER bytes:              30

Per-DSL adjustment for reporting:
  ADSL adjustment factor:  100
  ADSL2 adjustment factor: 100
  ADSL2+ adjustment factor:100
  VDSL adjustment factor:  100
  VDSL2 adjustment factor: 100
  SDSL adjustment factor:  100
  OTHER adjustment factor:  100

Keepalive Timer:         45 secs
State:                   WRITE_READY
Connect Time:            Fri May  2 12:08:49 2016
Session Time:            Fri May  2 12:18:52 2016
Routing Instance Time:   Fri May  2 12:18:53 2016
Keepalive Time:          Fri May  2 13:44:34 2016
Update Time:             Fri May  2 13:02:55 2016
```

show ancp neighbor

Syntax

show ancp neighbor
<brief | detail>
<ip-address *ip-address*
<system-name *mac-address*>

Release Information Command introduced in Junos OS Release 9.4.

Description Display information about all ANCP neighbors or the specified ANCP neighbor, regardless of operational state.

Options

brief | detail—(Optional) Display the specified level of detail.

ip-address ***ip-address*** —(Optional) Display information about the neighbor (access node) specified by the IP address.

system-name ***mac-address***—(Optional) Display information about the neighbor (access node) specified by the MAC address.

Required Privilege Level view

Related Documentation

- show ancp cos on page 1619
- show ancp subscriber on page 1637

List of Sample Output

show ancp neighbor on page 1627
show ancp neighbor detail on page 1628
show ancp neighbor ip-address on page 1629
show ancp neighbor system-name on page 1630

Output Fields Table 81 on page 1624 lists the output fields for the **show ancp neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 81: show ancp neighbor Output Fields

Field Name	Field Description	Level of Output
Version	Version of the ANCP implementation: • **0x31**—General Switch Management Protocol (GSMP) version 3, sub-version 1; ANCP version before *RFC 6320, Protocol for Access Node Control Mechanism in Broadband Networks*. • **0x32**—ANCP version 1, defined in *RFC 6320, Protocol for Access Node Control Mechanism in Broadband Networks*.	**brief detail** none
IP Address	IP address of the ANCP neighbor.	**brief detail** none

Copyright © 2018, Juniper Networks, Inc.

Table 81: show ancp neighbor Output Fields (continued)

Field Name	Field Description	Level of Output
PartId	Number that associates the ANCP message with a specific partition.	**brief** none
State	Operational state of the ANCP adjacency: • Configured—The neighbor has been configured, but has never been in the **Established** state. An asterisk (*) is prefixed to the neighbor entry for this state. • Establishing—Adjacency negotiations are in progress for the neighbor. An asterisk (*) is prefixed to the neighbor entry for this state. This state is rarely seen because the adjacency is established so quickly. • **Established**—Adjacency negotiations have succeeded for the neighbor and an ANCP session has been established. • Not Estblshed—Not Established; adjacency negotiations are ready to begin. Indicates that this neighbor previously had been in the **Established** state; that is, it has lost a previously established adjacency. An asterisk (*) is prefixed to the neighbor entry for this state.	All levels
Time	How long the adjacency has been up in one of the following formats: • *nwndnh*—number of weeks, days, and hours • *nd hh:mm:ss*—number of days, hours, minutes, and seconds	**brief detail** none
Subscriber Count	Number of subscribers associated with the ANCP neighbor (access local loop).	**brief** none
Capabilities	Negotiated ANCP capability: • **Topo**—Topology discovery. • **OAM**—Performance of local Operations Administration Maintenance (OAM) procedures on an access loop controlled by the router.	All levels
System Name	MAC address of the ANCP neighbor.	**detail**
TCP Port	TCP port on which ANCP messages are exchanged.	**detail**
System Instance	Number identifying the ANCP link instance from the edge device's perspective.	**detail**
Peer Instance	Number identifying the ANCP instance from the access node's perspective. This number is unique and changes when the node or link comes back up after going down.	**detail**
Timer	Adjacency timer value advertised by the ANCP peer in 100 ms increments; the interval between ANCP ACK messages. This value remains constant for the duration of an ANCP session.	**detail**
Partition Type	Number that identifies whether partitions are used and how the ID is negotiated: • **0**—No partition. • **1**—Fixed partition requested. • **2**—Fixed partition assigned.	**detail**

Copyright © 2018, Juniper Networks, Inc.

Table 81: show ancp neighbor Output Fields (continued)

Field Name	Field Description	Level of Output
Partition Flag	Number that specifies the type of partition requested: **1** (new adjacency) or **2** (recovered adjacency).	**detail**
Partition Identifier	Number that identifies a logical partition of an access node with which the ANCP agent has formed an adjacency. A value of zero indicates that the agent supports each neighbor on an IP address over a single TCP session with a partition ID of zero. This is the default support case. A nonzero value indicates that the agent supports each neighbor on an IP address over a single TCP session with a nonzero partition ID.	**detail**
Partition Adjacencies	Number of adjacencies that share the partition.	**detail**
Dead Timer	Remaining period that the edge device waits for adjacency packets from a neighbor before declaring the neighbor to be down. The maximum dead time value is three times the configured adjacency timer value. This field displays the current value based on the time that the last adjacency packet was received.	**detail**
Received Syn Count	Number of synchronization messages received from neighbors to maintain adjacencies.	**detail**
Received Synack Count	Number of synchronization acknowledgment messages received from neighbors in response to the node’s synchronization messages.	**detail**
Received Rstack Count	Number of messages received from neighbors indicating that the link to the neighbor needs to be reset.	**detail**
Received Ack Count	Number of acknowledgment messages periodically received from neighbors after an adjacency has been established.	**detail**
Received Port Up Count	Number of status messages received from neighbors indicating that a port has transitioned to the up state.	**detail**
Received Port Down Count	Number of status messages received from neighbors indicating that a port has transitioned to the down state.	**detail**
Received Generic Resp Count	Number of generic response messages received from neighbors.	**detail**
Received Adjacency Update Count	Number of adjacency update messages received from neighbors.	**detail**
Received OAM Count	Number of OAM responses received from neighbors in reply to **request** commands.	**detail**
Received Other Count	Number of all other ANCP message packets received from neighbors that do not fit into one of the other categories.	**detail**
Sent Syn Count	Number of synchronization messages sent to neighbors to maintain adjacencies.	**detail**

 Copyright © 2018, Juniper Networks, Inc.

Table 81: show ancp neighbor Output Fields (continued)

Field Name	Field Description	Level of Output
Sent Synack Count	Number of synchronization acknowledgment messages sent to neighbors in response to the their synchronization messages.	**detail**
Sent Rstack Count	Number of messages sent to neighbors indicating that the link to the neighbor needs to be reset.	**detail**
Sent Ack Count	Number of acknowledgment messages periodically sent to neighbors after an adjacency has been established.	**detail**
Sent Generic Resp Count	Number of generic response messages sent to neighbors.	**detail**
Sent OAM Count	Number of OAM **request** commands sent to neighbors.	**detail**
Max Discovery Limit Exceed Count	Number of times that the maximum number of discovery table entries accepted from the neighbor has been exceeded.	**detail**
Result Codes	Number of generic response messages sent to neighbors that include each of the following result codes: • **Invalid Request Message Count**—A properly formed request message violated the protocol because of timing (such as a race condition) or direction of transmission. • **Specified Port(s) Down Count**—One or more of the specified ports are down because of a state mismatch between the router and an ANCP control application. • **Out of Resources Count**—ANCP is out of resources, probably not related to the access lines. This result code is sent only by an access node. • **Request Msg Not Implemented Count**— • **Malformed Msg Count**—Message is malformed because it was corrupted in transit or there was an implementation error at either end of the connection. • **TLV Missing Count**—One or more mandatory TLVs was missing from a request. • **Invalid TLV Contents Count**—The contents of one or more TLVs in the request do not match its required specification. • **Non-Existent Port(s) Count**—One or more of the ports specified in a request do not exist, possibly because of a configuration mismatch between the access node and the router or AAA.	**detail**

Sample Output

show ancp neighbor

```
user@host> show ancp neighbor
  Version IP Address       PartID  State         Time        Subscriber
Capabilities
                                                             Count
  0x31    203.0.113.13         0       Established   11:24       2
Topo
  0x31    203.0.113.15         0       Not Estblshd  2:45        2
Topo
* 0x0     198.51.100.102       0       Establishing  0           0
```

```
* 0x0      192.0.2.0           0       Configured    0            0
* 0x0      192.0.2.1           0       Configured    0            0
```

show ancp neighbor detail

```
user@host> show ancp neighbor detail
  Neighbor Information
      Version                  : 0x31
      IP Address               : 192.0.2.85
      System Name              : 00:00:5e:00:53:01
          Up Time                            : 26
          TCP Port                           : 32666
          State                              : Established
          Subscriber Count                   : 4
          Capabilities                       : Topo
          System Instance                    : 2
          Peer Instance                      : 20
          Adjacency Timer (in 100ms)         : 100
          Peer Adjacency Timer (in 100ms)    : 100
          Partition Type                     : 0
          Partition Flag                     : 1
          Partition Identifier               : 0
          Partition Adjacencies              : 0
          Dead Timer                         : 23
          Received Syn Count                 : 1
          Received Synack Count              : 1
          Received Rstack Count              : 0
          Received Ack Count                 : 4
          Received Port Up Count             : 10
          Received Port Down Count           : 0
          Received Generic Resp Count        : 0
          Received Adjacency Update Count    : 0
          Received OAM Count                 : 0
          Received Other Count               : 0
          Sent Syn Count                     : 1
          Sent Synack Count                  : 2
          Sent Rstack Count                  : 0
          Sent Ack Count                     : 3
          Sent Generic Resp Count            : 0
          Sent OAM Count                     : 0
          Max Discovery Limit Exceed Count : 0
      Result Codes:                          Received              Sent
          Invalid Request Message Count    : 0                     0
          Specified Port(s) Down Count     : 0                     0
          Out of Resources Count           : 0                     0
          Request Msg Not Implemented Count: 0                     0
          Malformed Msg Count              : 0                     0
          TLV Missing Count                : 0                     0
          Invalid TLV Contents Count       : 0                     0
          Non-Existent Port(s) Count       : 0                     0


      Version                  : 0x32
      IP Address               : 192.168.9.1
      System Name              : 00:00:5e:00:53:02
          Up Time                            : 36
          TCP Port                           : 61408
          State                              : Not Established
          Subscriber Count                   : 1
          Capabilities                       : Topology Discovery
          System Instance                    : 12
```

 Copyright © 2018, Juniper Networks, Inc.

```
        Peer Instance                    : 1
        Adjacency Timer (in 100ms)       : 50
        Peer Adjacency Timer (in 100ms)  : 100
        Partition Type                   : 0
        Partition Flag                   : 1
        Partition Identifier             : 0
        Partition Adjacencies            : 0
        Dead Timer                       : 23
        Received Syn Count               : 24
        Received Synack Count            : 20
        Received Rstack Count            : 2
        Received Ack Count               : 9
        Received Port Up Count           : 5
        Received Port Down Count         : 0
        Received Generic Resp Count      : 0
        Received Adjacency Update Count  : 0
        Received OAM Responses Count     : 2
        Received Other Count             : 0
        Sent Syn Count                   : 20
        Sent Synack Count                : 24
        Sent Rstack Count                : 1
        Sent Generic Resp Count          : 0
        Sent Ack Count                   : 9
        Sent OAM Requests Count          : 4
        Max Discovery Limit Exceed Count : 0
    Result Codes:                          Received              Sent
        Invalid Request Message Count    : 0                     0
        Specified Port(s) Down Count     : 0                     0
        Out of Resources Count           : 0                     0
        Request Msg Not Implemented Count: 0                     0
        Malformed Msg Count              : 0                     0
        TLV Missing Count                : 0                     0
        Invalid TLV Contents Count       : 0                     0
        Non-Existent Port(s) Count       : 0                     0
```

show ancp neighbor ip-address

user@host> **show ancp neighbor ip-address 192.0.2.85**

```
Neighbor Information
      Version                 : 0x32
      IP Address    : 192.0.2.85
      System Name   : 00:00:5e:00:53:ba
          Up Time                          : 26
          TCP Port                         : 32666
          State                            : Established
          Subscriber Count                 : 4
          Capabilities                     : Topo
          System Instance                  : 2
          Peer Instance                    : 20
          Adjacency Timer (in 100ms)       : 100
          Peer Adjacency Timer (in 100ms)  : 100
          Partition Type                   : 0
          Partition Flag                   : 1
          Partition Identifier             : 0
          Partition Adjacencies            : 0
          Dead Timer                       : 23
          Received Syn Count               : 1
          Received Synack Count            : 1
          Received Rstack Count            : 0
          Received Ack Count               : 4
```

```
        Received Port Up Count          : 10
        Received Port Down Count        : 0
        Received Generic Resp Count     : 0
        Received Adjacency Update Count : 0
        Received OAM Count              : 0
        Received Other Count            : 0
        Sent Syn Count                  : 1
        Sent Synack Count               : 2
        Sent Rstack Count               : 0
        Sent Ack Count                  : 3
        Sent Generic Resp Count         : 0
        Sent OAM Count                  : 0
        Max Discovery Limit Exceed Count : 0
    Result Codes:                         Received              Sent
        Invalid Request Message Count   : 0                     0
        Specified Port(s) Down Count    : 0                     0
        Out of Resources Count          : 0                     0
        Request Msg Not Implemented Count: 0                    0
        Malformed Msg Count             : 0                     0
        TLV Missing Count               : 0                     0
        Invalid TLV Contents Count      : 0                     0
        Non-Existent Port(s) Count      : 0                     0
```

show ancp neighbor system-name

```
user@host> show ancp neighbor 00:00:5e:00:53:ba detail

Neighbor Information
   Version                 : 0x31
   IP Address              : 203.0.113.101
   System Name             : 00:00:5e:00:53:ba
     Up Time                              : 19
     TCP Port                             : 1028
     State                                : Established
     Subscriber Count                     : 2
     Capabilities                         : Topology Discovery, OAM
     System Instance                      : 1
     Peer Instance                        : 10
     Adjacency Timer (in 100ms)           : 100
     Peer Adjacency Timer (in 100ms)      : 250
     Partition Type                       : 0
     Partition Flag                       : 1
     Partition Identifier                 : 0
     Partition Adjacencies                : 0
     Dead Timer                           : 55
     Received Syn Count                   : 1

     Received Synack Count                : 1
     Received Rstack Count                : 0
     Received Ack Count                   : 1
     Received Port Up Count               : 34
     Received Port Down Count             : 0
     Received Generic Resp Count          : 0
     Received Adjacency Update Count      : 0
     Received OAM Responses Count         : 2
     Received Other Count                 : 0
     Sent Syn Count                       : 1
     Sent Synack Count                    : 1
     Sent Rstack Count                    : 0
     Sent Ack Count                       : 3
     Sent Generic Resp Count              : 0
```

 Copyright © 2018, Juniper Networks, Inc.

```
  Sent OAM Requests Count            : 4
  Max Discovery Limit Exceed Count   : 3
Result Codes:                        Received             Sent
  Invalid Request Message Count      : 0                    0
  Specified Port(s) Down Count       : 0                    0
  Out of Resources Count             : 0                    0
  Request Msg Not Implemented Count  : 0                    0
  Malformed Msg Count                : 0                    0
  TLV Missing Count                  : 0                    0
  Invalid TLV Contents Count         : 0                    0
  Non-Existent Port(s) Count         : 0                    0
```

Copyright © 2018, Juniper Networks, Inc.

show ancp statistics

Syntax

show ancp statistics
<ip-address *ip-address*>
<system-name *mac-address*>

Release Information

Command introduced in Junos OS Release 13.3.

Description

Display statistics for all ANCP neighbors (access nodes) or the specified ANCP neighbor.

Options

none—Display statistics for all ANCP neighbors, including global statistics not show for individual neighbors.

ip-address *ip-address*—(Optional) Display statistics for only the neighbor with the specified IP address.

system-name *mac-address*—(Optional) Display statistics for only the neighbor with the specified MAC address.

Required Privilege Level

view

Related Documentation

- show ancp cos on page 1619
- show ancp neighbor on page 1624
- show ancp subscriber on page 1637

List of Sample Output

show ancp statistics on page 1634
show ancp statistics ip-address on page 1635
show ancp statistics system-name on page 1635

Output Fields

Table 82 on page 1632 lists the output fields for the **show ancp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 82: show ancp statistics Output Fields

Field Name	Field Description
Number of neighbors	Total count of ANCP neighbors.
Number of subscribers	Total count of ANCP subscribers.
Accept Count	Number of neighbor TCP/IP sessions accepted on listener socket.

 Copyright © 2018, Juniper Networks, Inc.

Table 82: show ancp statistics Output Fields (continued)

Field Name	Field Description
Accept Fail Count	Number of neighbor TCP/IP sessions that failed due to one of the following causes: session already exists, maximum number of ANCP connections exceeded, creation of session or neighbor failed, or protocol start failed.
No Config Accept Deny Count	Number of neighbor TCP/IP sessions that failed because the neighbor was not configured.
Received Syn Count	Number of synchronization messages received from neighbors to maintain adjacencies.
Received Synack Count	Number of synchronization acknowledgment messages received from neighbors in response to the node's synchronization messages.
Received Rstack Count	Number of messages received from neighbors indicating that the link to the neighbor needs to be reset.
Received Ack Count	Number of acknowledgment messages periodically received from neighbors after an adjacency has been established.
Received Port Up Count	Number of status messages received from neighbors indicating that a port has transitioned to the up state.
Received Port Down Count	Number of status messages received from neighbors indicating that a port has transitioned to the down state.
Received Generic Resp Count	Number of generic response messages received from neighbors.
Received Adjacency Update Count	Number of adjacency update messages received from neighbors.
Received OAM Count	Number of OAM responses received from neighbors in reply to **request** commands.
Received Other Count	Number of all other ANCP message packets received from neighbors that do not fit into one of the other categories.
Sent Syn Count	Number of synchronization messages sent to neighbors to maintain adjacencies.
Sent Synack Count	Number of synchronization acknowledgment messages sent to neighbors in response to the their synchronization messages.
Sent Rstack Count	Number of messages sent to neighbors indicating that the link to the neighbor needs to be reset.
Sent Ack Count	Number of acknowledgment messages periodically sent to neighbors after an adjacency has been established.
Sent Generic Resp Count	Number of generic response messages sent to neighbors.

Copyright © 2018, Juniper Networks, Inc.

Table 82: show ancp statistics Output Fields (continued)

Field Name	Field Description
Sent OAM Count	Number of OAM **request** commands sent to neighbors.
Result Codes	Number of generic response messages sent to neighbors that include each of the following result codes: • **Invalid Request Message Count**—A properly formed request messages violated the protocol because of timing (such as a race condition) or direction of transmission. • **Specified Port(s) Down Count**—One or more of the specified ports are down because of a state mismatch between the router and an ANCP control application. • **Out of Resources Count**—the ANCP agent is out of resources, probably not related to the access lines. This result code is sent only by an access node. • **Request Msg Not Implemented Count**— • **Malformed Msg Count**—Message is malformed because it was corrupted in transit or there was an implementation error at either end of the connection. • **TLV Missing Count**—One or more mandatory TLVs was missing from a request. • **Invalid TLV Contents Count**—The contents of one or more TLVs in the request do not match its required specification. • **Non-Existent Port(s) Count**—One or more of the ports specified in a request do not exist, possibly because of a configuration mismatch between the access node and the router or AAA.

Sample Output

show ancp statistics

```
user@host> show ancp statistics
  Statistics
          Number of neighbors             : 4
          Number of subscribers           : 6
          Accept Count                    : 0
          Accept Fail Count               : 0
          No Config Accept Deny Count     : 0
          Received Syn Count              : 2
          Received Synack Count           : 1
          Received Rstack Count           : 0
          Received Ack Count              : 8
          Received Port Up Count          : 7
          Received Port Down Count        : 0
          Received Generic Resp Count     : 0
          Received Adjacency Update Count : 0
          Received OAM Count              : 0
          Received Other Count            : 0
          Sent Syn Count                  : 1
          Sent Synack Count               : 1
          Sent Rstack Count               : 0
          Sent Ack Count                  : 17
          Sent Generic Resp Count         : 0
          Sent OAM Count                  : 4
      Result Codes:                         Received               Sent
          Invalid Request Message Count    : 0                     0
          Specified Port(s) Down Count     : 0                     0
```

 Copyright © 2018, Juniper Networks, Inc.

```
        Out of Resources Count          : 0                     0
        Request Msg Not Implemented Count: 0                    0
        Malformed Msg Count             : 0                     0
        TLV Missing Count               : 0                     0
        Invalid TLV Contents Count      : 0                     0
        Non-Existent Port(s) Count      : 0                     0
```

show ancp statistics ip-address

```
user@host> show ancp statistics ip-address 203.0.113.1
  Statistics
        Received Syn Count              : 2
        Received Synack Count           : 1
        Received Rstack Count           : 0
        Received Ack Count              : 8
        Received Port Up Count          : 7
        Received Port Down Count        : 0
        Received Generic Resp Count     : 0
        Received Adjacency Update Count : 0
        Received OAM Count              : 0
        Received Other Count            : 0
        Sent Syn Count                  : 1
        Sent Synack Count               : 1
        Sent Rstack Count               : 0
        Sent Ack Count                  : 17
        Sent Generic Resp Count         : 0
        Sent OAM Count                  : 4
    Result Codes:                         Received              Sent
        Invalid Request Message Count    : 0                     0
        Specified Port(s) Down Count     : 0                     0
        Out of Resources Count           : 0                     0
        Request Msg Not Implemented Count: 0                     0
        Malformed Msg Count              : 0                     0
        TLV Missing Count                : 0                     0
        Invalid TLV Contents Count       : 0                     0
        Non-Existent Port(s) Count       : 0                     0
```

show ancp statistics system-name

```
user@host> show ancp statistics system-name 00:00:5E:00:53:02
  Statistics
        Received Syn Count              : 2
        Received Synack Count           : 1
        Received Rstack Count           : 0
        Received Ack Count              : 8
        Received Port Up Count          : 7
        Received Port Down Count        : 0
        Received Generic Resp Count     : 0
        Received Adjacency Update Count : 0
        Received OAM Count              : 0
        Received Other Count            : 0
        Sent Syn Count                  : 1
        Sent Synack Count               : 1
        Sent Rstack Count               : 0
        Sent Ack Count                  : 17
        Sent Generic Resp Count         : 0
        Sent OAM Count                  : 4
    Result Codes:                         Received              Sent
```

Copyright © 2018, Juniper Networks, Inc.

```
Invalid Request Message Count    : 0                     0
Specified Port(s) Down Count     : 0                     0
Out of Resources Count           : 0                     0
Request Msg Not Implemented Count: 0                     0
Malformed Msg Count              : 0                     0
TLV Missing Count                : 0                     0
Invalid TLV Contents Count       : 0                     0
Non-Existent Port(s) Count       : 0                     0
```

 Copyright © 2018, Juniper Networks, Inc.

show ancp subscriber

Syntax

show ancp subscriber
<brief | detail>
<identifier *identifier*>
<ip-address *ip-address*>
<system-name *mac-address*>

Release Information

Command introduced in Junos OS Release 9.4.
neighbor option replaced with **ip-address** in Junos OS Release 16.1.
system-name option introduced in Junos OS Release 16.1.

Description

Display information about active subscribers regardless of the subscriber's operational state, for all subscribers (local access loops), the subscriber associated with the access line specified by an ACI, or the subscriber associated with the specified ANCP neighbor (access node).

After an ancpd restart, this command displays orphaned entries (marked with an **o**) for subscriber sessions that were established before the restart but which have not yet been reestablished. As sessions are reestablished, the number of orphaned entries displayed by the command decreases. The number reaches zero when all sessions are reestablished or when the orphaned-interface timer expires.

Options

none—Display information about all subscribers.

brief | detail—(Optional) Display the specified level of detail.

identifier *identifier*—(Optional) Display information about the subscriber associated with the access line (ACI) specified by the access identifier.

ip-address *ip-address* —(Optional) Display information about the subscribers connected to the access node specified by the IP address.

system-name *mac-address*—(Optional) Display information about the subscribers connected to the access node specified by the MAC address.

Required Privilege Level

view

Related Documentation

- clear ancp subscriber on page 1614
- show ancp cos on page 1619
- show ancp neighbor on page 1624
- show ancp statistics on page 1632

List of Sample Output

show ancp subscriber on page 1640
show ancp subscriber (After ancpd Restart) on page 1641

show ancp subscriber brief on page 1641
show ancp subscriber detail on page 1641
show ancp subscriber identifier identifier-string detail on page 1642

Output Fields Table 83 on page 1638 lists the output fields for the **show ancp subscriber** command. Output fields are listed in the approximate order in which they appear.

Table 83: show ancp subscriber Output Fields

Field Name	Field Description	Level of Output
Loop Identifier	Access loop identifier as sent by the access node and configured to map the subscriber to an interface. An asterisk (*) indicates that the information might be stale due to receiving a Port Down message with a DSL Line State of Idle. Two asterisks (**) indicate that the neighbor associated with the subscriber has lost its adjacency. In this case, the DSL Line State might be Established. An **o** indicates that the entry is for an orphaned interface and represents a previously established subscriber session that has not been reestablished after an ancpd restart. The number of orphaned entries decreases as the ANCP neighbors reestablish adjacencies and the protocol subscriber sessions are reestablished. The command output indicates this by removing the **o** marker. Eventually the number of orphaned entries reaches zero, because either all the adjacencies and subscriber sessions have been reestablished or any remaining orphaned entries are removed when the orphaned-interface timer expires.	**brief** none
DSL Line State	State of the DSL line: **Idle**, **Showtime**, or **Silent**.	**brief detail**
Access Type	Type of access line employed by the access node: **ADSL1**, **ADSL2**, **ADSL2+**, **VDSL1**, **VDSL2**, **SDSL**, **G.fast**, **VDSL2 Annex Q**, **SDSL bonded**, **VDSL2 bonded**, **G.fast bonded VDSL2 Annex Q bonded** or **OTHER**.	**brief detail** none
Interface	Name of the interface set or logical interface.	**brief detail** none
Rate Kbps	Actual downstream data rate for this local loop.	**brief** none
Neighbor	IP address of ANCP neighbor (access node).	**brief** none

 Copyright © 2018, Juniper Networks, Inc.

Table 83: show ancp subscriber Output Fields (continued)

Field Name	Field Description	Level of Output
Access Loop Circuit Identifier	Access loop circuit identifier as sent by the access node and configured to map the subscriber to an interface. An asterisk (*) indicates that the information might be stale due to receiving a Port Down message with a DSL Line State of Idle. Two asterisks (**) indicate that the neighbor associated with the subscriber has lost its adjacency. In this case, the DSL Line State might be Established.	**detail**
Neighbor IP Address	IP address of the ANCP neighbor (access node).	**detail**
Aggregate Circuit Identifier Binary	Binary identifier for the VLAN circuit ID.	**detail**
Tech Type	Type of technology employed by the subscriber. Currently Junos OS supports **DSL** technology type only.	**detail**
DSL Line Data Link	Data link protocol employed on the access loop: **AAL5** or **Ethernet**.	**detail**
DSL Line Encapsulation	Encapsulation type on the access loop, for Ethernet only: • **0**—NA, type not conveyed • **1**—Untagged Ethernet • **2**—Single-tagged Ethernet	**detail**
DSL Line Encapsulation Payload	Payload carried across the access loop: • **0**—NA, type not conveyed • **1**—PPPoA LLC • **2**—PPPoA null • **3**—IPoA LLC • **4**—IPoA null • **5**—Ethernet over AAL5 LLC with FCS • **6**—Ethernet over AAL5 LLC without FCS • **7**—Ethernet over AAL5 null with FCS • **8**—Ethernet over AAL5 null without FCS	**detail**
Interface Type	Type of interface employed for subscriber traffic: **ifl** for a single VLAN or **interface-set** for a configured group of VLANs.	**detail**
Actual Net Data Upstream	Actual upstream data rate for this local loop, in Kbps.	**detail**
Actual Net Data Downstream	Actual downstream data rate for this local loop, in Kbps.	**detail**

Copyright © 2018, Juniper Networks, Inc.

Table 83: show ancp subscriber Output Fields (continued)

Field Name	Field Description	Level of Output
Minimum Net Data Upstream	Minimum upstream data rate desired by the operator for this local loop, in Kbps.	**detail**
Minimum Net Data Downstream	Minimum downstream data rate desired by the operator for this local loop, in Kbps.	**detail**
Maximum Net Data Upstream	Maximum upstream data rate desired by the operator for this local loop, in Kbps.	**detail**
Maximum Net Data Downstream	Maximum downstream data rate desired by the operator for this local loop, in Kbps.	**detail**
Attainable Net Data Upstream	Maximum attainable upstream data rate for this local loop, in Kbps.	**detail**
Attainable Net Data Downstream	Maximum attainable downstream data rate for this local loop, in Kbps.	**detail**
Minimum Low Power Data Downstream	Minimum downstream data rate desired by the operator for this local loop in low power state, in Kbps.	**detail**
Minimum Low Power Data Upstream	Minimum upstream data rate desired by the operator for this local loop in low power state, in Kbps.	**detail**
Maximum Interleave Delay Downstream	Maximum interleaving delay for downstream data, in milliseconds.	**detail**
Maximum Interleave Delay Upstream	Maximum interleaving delay for upstream data, in milliseconds.	**detail**
Actual Interleave Delay Downstream	Actual interleaving delay for downstream data, in milliseconds.	**detail**
Actual Interleave Delay Upstream	Actual interleaving delay for upstream data, in milliseconds.	**detail**

Sample Output

show ancp subscriber

```
user@host> show ancp subscriber
  Loop Identifier        DSL Line  Tech Type     Access Type   Interface
Rate   Neighbor
                         State                              Kbps
**circuit 101            Idle      DSL       ADSL1    ----             32
 203.0.113.13
**circuit 102            Idle      DSL       ADSL1    ----             32
 203.0.113.13
```

 Copyright © 2018, Juniper Networks, Inc.

```
 circuit 301          Showtime  DSL       ADSL1    ----              32
203.0.113.15
 circuit 302          Showtime  DSL       ADSL1    ----              32
203.0.113.15
```

show ancp subscriber (After ancpd Restart)

```
user@host> show ancp subscriber
  Loop Identifier       DSL Line  Tech Type     Access Type  Interface
Rate     Neighbor
                        State                                Kbps
o circuit 201           Showtime  DSL        ADSL1    ----              222222

o circuit 202           Showtime  DSL        ADSL1    ----              222222
```

show ancp subscriber brief

```
user@host> show ancp subscriber brief

  Loop Identifier         Type      Interface        Rate         Neighbor
                                                     Kbps
  port-1-10              VDSL2     set-ge-10410      64        203.0.113.102

  port-1-11              VDSL2     set-ge-10411      64        203.0.113.111

  port-2-10              VDSL2     ge-1/0/4.12       64        203.0.113.112
  port-2-11              VDSL2     ge-1/0/4.13       64        203.0.113.113
```

show ancp subscriber detail

```
user@host> show ancp subscriber detail
 Subscriber Information
*  Access Loop Circuit Identifier : circuit 101
     Neighbor IP Address                : 203.0.113.13
     Aggregate Circuit Identifier Binary  : 0/0
     Tech Type                : DSL
     Access Type                        : ADSL1
     DSL Line State                     : Idle
     DSL Line Data Link                 : Data link 2
     DSL Line Encapsulation             : N/A
     DSL Line Encapsulation Payload     : N/A
     Interface Type                     : N/A
     Interface                          : ----
     Actual Net Data Upstream           : 32
     Actual Net Data Downstream         : 32
     Minimum Net Data Upstream          : 0
     Minimum Net Data Downstream        : 0
     Maximum Net Data Upstream          : 0
     Maximum Net Data Downstream        : 0
     Attainable Net Data Upstream       : 1024
     Attainable Net Data Downstream     : 8192
     Minimum Low Power Data Downstream  : 32
     Minimum Low Power Data Upstream    : 32
     Maximum Interleave Delay Downstream  : 20
     Maximum Interleave Delay Upstream  : 20
     Actual Interleave Delay Downstream  : 20
     Actual Interleave Delay Upstream   : 20
*  Access Loop Circuit Identifier: circuit 102
```

```
    Neighbor IP Address                  : 213.0.113.13
    Aggregate Circuit Identifier Binary  : 0/0
    Tech Type                  : DSL
    Access Type                          : ADSL1
    DSL Line State                       : Idle
    DSL Line Data Link                   : Data link 2
    DSL Line Encapsulation               : N/A
    DSL Line Encapsulation Payload       : N/A
    Interface Type                       : N/A
    Interface                            : ----
    Actual Net Data Upstream             : 32
    Actual Net Data Downstream           : 32
    Minimum Net Data Upstream            : 0
    Minimum Net Data Downstream          : 0
    Maximum Net Data Upstream            : 0
    Maximum Net Data Downstream          : 0
    Attainable Net Data Upstream         : 1024
    Attainable Net Data Downstream       : 8192
    Minimum Low Power Data Downstream    : 32
    Minimum Low Power Data Upstream      : 32
    Maximum Interleave Delay Downstream  : 20
    Maximum Interleave Delay Upstream    : 20
    Actual Interleave Delay Downstream   : 20
    Actual Interleave Delay Upstream     : 20
...
```

show ancp subscriber identifier identifier-string detail

user@host> **show ancp subscriber identifier port-1-11 detail**

```
 Access Loop Identifier : port-1-11
     Neighbor IP Address                  : 203.0.113.112
     Aggregate Circuit Identifier Binary  : 0/0
     DSL Type                             :  DSL  0
     Interface Type                       : interface-set
     Interface                            : set-ge-10411
     DSL Line State                       : Show Time
     Actual Net Data Upstream             : 64
     Actual Net Data Downstream           : 64
     DSL Line Data Link                   : AAL5
     DSL Line Encapsulation               : N/A
     DSL Line Encapsulation Payload       : N/A
     Minimum Net Data Upstream            : 64
     Minimum Net Data Downstream          : 64
     Maximum Net Data Upstream            : 64
     Maximum Net Data Downstream          : 64
     Attainable Net Data Upstream         : 64
     Attainable Net Data Downstream       : 64
     Minimum Low Power Data Downstream    : 64
     Minimum Low Power Data Upstream      : 64
     Maximum Interleave Delay Downstream  : 50
     Maximum Interleave Delay Upstream    : 50
     Actual Interleave Delay Downstream   : 50
     Actual Interleave Delay Upstream     : 50
```


Copyright © 2018, Juniper Networks, Inc.

show ancp summary

Syntax show ancp summary

Release Information Command introduced in Junos OS Release 13.1.

Description Display a summary of the counts and states for all ANCP neighbors and subscribers.

Options This command has no options.

Required Privilege Level view

Related Documentation

- show ancp neighbor on page 1624
- show ancp summary neighbor on page 1645
- show ancp subscriber on page 1637
- show ancp summary subscriber on page 1647

List of Sample Output show ancp summary on page 1644

Output Fields Table 84 on page 1643 lists the output fields for the **show ancp summary** command. Output fields are listed in the approximate order in which they appear.

Table 84: show ancp summary Output Fields

Field Name	Field Description
Configured	Number of ANCP neighbors in the **Configured** state; that is, that have been configured but never established.
Establishing	Number of ANCP neighbors in the **Establishing** state; that is, where negotiations are in progress.
Established	Number of ANCP neighbors in the **Established** state; that is, where negotiations have succeeded and the ANCP session has been established.
Not Estblshd	Number of ANCP neighbors in the **Not Estblshd** state; that is, that have lost a previously established adjacency and are ready to begin negotiations.
Total	Total number of ANCP neighbors; sum of neighbors in the **Configured**, **Establishing**, **Established**, and **Not Estblshd** states.
Showtime	Number of DSL lines in **Showtime** state.
Idle	Number of DSL lines in **Idle** state.
Silent	Number of DSL lines in **Silent** state.

Copyright © 2018, Juniper Networks, Inc.

Table 84: show ancp summary Output Fields (continued)

Field Name	Field Description
Unknown	Number of DSL lines where the state is not **Showtime**, **Idle**, or **Silent**.
Total	Total number of DSL lines (ANCP subscribers); sum of DSL lines in the **Showtime**, **Idle**, **Silent**, and **Unknown** states.

Sample Output

show ancp summary

```
user@host> show ancp summary

Neighbors Summary:
  Configured  Establishing  Established  Not Established     Total
  ----------  ------------  -----------  ---------------  ----------
          22             0            2                0          24

Subscribers Summary:
  Showtime       Idle      Silent      Unknown      Total
  ---------- ---------- ----------- ---------- ----------
           4          0           0          0          4
```


Copyright © 2018, Juniper Networks, Inc.

show ancp summary neighbor

Syntax

show ancp summary neighbor
<**ip-address** *ip-address* | **system-name** *mac-address*>

Release Information

Command introduced in Junos OS Release 13.1.

Description

Display a summary of the counts and states for all ANCP neighbors and of the neighbor's subscribers when you specify a particular neighbor.

Options

ip-address *ip-address*—(Optional) IP address of the ANCP neighbor (access node).

system-name *mac-address*—(Optional) MAC address of the ANCP neighbor (access node).

Required Privilege Level

view

Related Documentation

- show ancp summary on page 1643
- show ancp subscriber on page 1637
- show ancp summary subscriber on page 1647

List of Sample Output

show ancp summary neighbor on page 1646
show ancp summary neighbor (IP Address) on page 1646
show ancp summary neighbor (MAC Address) on page 1646

Output Fields

Table 85 on page 1645 lists the output fields for the **show ancp summary** command. Output fields are listed in the approximate order in which they appear.

Table 85: show ancp summary neighbor Output Fields

Field Name	Field Description
Configured	Number of ANCP neighbors in the **Configured** state; that is, that have been configured but never established.
Establishing	Number of ANCP neighbors in the **Establishing** state; that is, where negotiations are in progress.
Established	Number of ANCP neighbors in the **Established** state; that is, where negotiations have succeeded and the ANCP session has been established.
Not Estblshd	Number of ANCP neighbors in the **Not Estblshd** state; that is, that have lost a previously established adjacency and are ready to begin negotiations.
Total	Total number of ANCP neighbors; sum of neighbors in the **Configured**, **Establishing**, **Established**, and **Not Estblshd** states.

Table 85: show ancp summary neighbor Output Fields (continued)

Field Name	Field Description
Showtime	Number of DSL lines for the neighbor in **Showtime** state.
Idle	Number of DSL lines for the neighbor in **Idle** state.
Silent	Number of DSL lines for the neighbor in **Silent** state.
Unknown	Number of DSL lines for the neighbor where the state is not **Showtime**, **Idle**, or **Silent**.
Total	Total number of DSL lines (ANCP subscribers); sum of DSL lines in the **Showtime**, **Idle**, **Silent**, and **Unknown** states.

Sample Output

show ancp summary neighbor

```
user@host> show ancp summary neighbor

Neighbors Summary:
  Configured  Establishing  Established  Not Established     Total
  ----------  ------------  -----------  ---------------  ----------
          22             0            2                0          24
```

show ancp summary neighbor (IP Address)

```
user@host> show ancp summary neighbor ip-address 192.168.10.1

Neighbor Summary:192.168.10.1 status Established

Subscribers Summary:
  Show Time      Idle       Silent      Unknown      Total
  ---------- ---------- ----------- ---------- ----------
           6          0           0          0          6
```

show ancp summary neighbor (MAC Address)

```
user@host> show ancp summary neighbor system-name 00:00:5E:00:53:02

Neighbor Summary:00:00:5E:00:53:02 status Established

Subscribers Summary:
  Show Time      Idle       Silent      Unknown      Total
  ---------- ---------- ----------- ---------- ----------
           5          1           2          0          8
```

 Copyright © 2018, Juniper Networks, Inc.

show ancp summary subscriber

Syntax show ancp summary subscriber

Release Information Command introduced in Junos OS Release 13.1.

Description Display a summary of the counts and states for all ANCP subscribers.

Options This command has no options.

Required Privilege Level view

Related Documentation

- show ancp summary on page 1643
- show ancp neighbor on page 1624
- show ancp summary neighbor on page 1645

List of Sample Output show ancp summary subscriber on page 1647

Output Fields Table 86 on page 1647 lists the output fields for the **show ancp summary subscriber** command. Output fields are listed in the approximate order in which they appear.

Table 86: show ancp summary subscriber Output Fields

Field Name	Field Description
Showtime	Number of DSL lines in **Showtime** state.
Idle	Number of DSL lines in **Idle** state.
Silent	Number of DSL lines in **Silent** state.
Unknown	Number of DSL lines where the state is not **Showtime**, **Idle**, or **Silent**.
Total	Total number of DSL lines (ANCP subscribers); sum of DSL lines in the **Showtime**, **Idle**, **Silent**, and **Unknown** states.

Sample Output

show ancp summary subscriber

```
user@host> show ancp summary subscriber

Subscribers Summary:
  Show Time     Idle      Silent     Unknown     Total
  ---------- ---------- ----------- ---------- ----------
           8          1           0          1         10
```

 Copyright © 2018, Juniper Networks, Inc.

CHAPTER 16

BFD Operational Commands

clear bfd adaptation

Syntax

```
clear bfd adaptation
<all>
<address session-address>
<discriminator discr-number>
```

Release Information Command introduced before Junos OS Release 7.4.

Description Clear adaptation for Bidirectional Forwarding Detection (BFD) sessions. BFD is a simple hello mechanism that detects failures in a network. Configured BFD interval timers can change, adapting to network situations. Use this command to return BFD interval timers to their configured values.

The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.

Options **all**—Clear adaptation for all BFD sessions.

address ***session-address***—(Optional) Clear adaptation for all BFD sessions matching the specified address.

discriminator ***discr-number***—(Optional) Clear adaptation for the local BFD session matching the specified discriminator.

Additional Information For more information, see the description of the **bfd-liveness-detection** configuration statement in the *Junos Routing Protocols Configuration Guide*.

Required Privilege Level clear

Related Documentation

- show bfd session on page 1652

List of Sample Output clear bfd adaptation on page 1650

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear bfd adaptation

```
user@host> clear bfd adaptation
```

 Copyright © 2018, Juniper Networks, Inc.

clear bfd session

List of Syntax

Syntax on page 1651
Syntax (EX Series Switch and QFX Series) on page 1651

Syntax

clear bfd session
<all>
<address *session-address*>
<discriminator *discr-number*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and QFX Series)

clear bfd session
<all>
<address *session-address*>
<discriminator *discr-number*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.1 for the QFX Series.

Description

Drop one or more Bidirectional Forwarding Detection (BFD) sessions.

Options

all—Drop all BFD sessions.

address *session-address*—(Optional) Drop all BFD sessions matching the specified address.

discriminator *discr-number*—(Optional) Drop the local BFD session matching the specified discriminator.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Related Documentation

- show bfd session on page 1652

List of Sample Output

clear bfd session all on page 1651

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear bfd session all

```
user@host> clear bfd session all
```

show bfd session

List of Syntax

Syntax on page 1652
Syntax (EX Series Switch and QFX Series) on page 1652

Syntax

show bfd session
<brief | detail | extensive | summary>
<address *address*>
<client rsvp-oam (brief | detail | extensive | summary) | vpls-oam (brief | detail | extensive |
instance *instance-name* | summary)>
<discriminator *discriminator*>
<logical-system (all | *logical-system-name*)>
<prefix *address*>
<subscriber (address *destination-address* | discriminator *discriminator* | extensive)>

Syntax (EX Series Switch and QFX Series)

show bfd session
<brief | detail | extensive | summary>
<address *address*>
<client rsvp-oam (brief | detail | extensive | summary) | vpls-oam (brief | detail | extensive |
instance *instance-name* | summary)>
<discriminator *discriminator*>
<prefix *address*>

Release Information

Command introduced before Junos OS Release 7.4.
Options **discriminator** and **address** introduced in Junos OS Release 8.2.
Option **prefix** introduced in Junos OS Release 9.0.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Option **client** introduced in Junos OS Release 12.3R3.
Option **subscriber** introduced in Junos OS Release 15.1 for the MX Series.

Description

Display information about active Bidirectional Forwarding Detection (BFD) sessions.

Options

none—(Same as **brief**) Display information about active BFD sessions.

brief | detail | extensive | summary—(Optional) Display the specified level of output.

address *address*—(Optional) Display information about the BFD session for the specified neighbor address.

client rsvp-oam
(brief | detail | extensive | summary)
| vpls-oam
(brief | detail | extensive | instance *instance-name* | summary)—(Optional) Display information about RSVP-OAM or VPLS-OAM BFD sessions in the specified level of output. For VPLS-OAM, display the specified level of output or display information about all of the BFD sessions for the specified VPLS routing instance.

discriminator *discriminator*—(Optional) Display information about the BFD session using the specified local discriminator.

 Copyright © 2018, Juniper Networks, Inc.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

<subscriber (address *destination-address* | discriminator *discriminator* | extensive)>—(Optional) Display information about all BFD sessions for subscribers, or for a single BFD subscriber session with a particular destination address, or with a particular denominator.

Required Privilege Level

view

Related Documentation

- clear bfd session on page 1651
- *Understanding BFD for Static Routes for Faster Network Failure Detection*
- *Example: Configuring BFD for Static Routes for Faster Network Failure Detection*
- *Understanding BFD for OSPF*
- *Example: Configuring BFD for OSPF*
- *Understanding BFD for BGP*
- *Example: Configuring BFD on Internal BGP Peer Sessions*
- *Understanding Bidirectional Forwarding Detection Authentication for PIM*
- *Configuring BFD for PIM*
- *Understanding BFD for IS-IS*

List of Sample Output

Output Fields

Table 87 on page 1654 describes the output fields for the **show bfd session** command. Output fields are listed in the approximate order in which they appear.

Table 87: show bfd session Output Fields

Field Name	Field Description	Level of Output
Address	Address on which the BFD session is active.	**brief detail extensive** none
State	State of the BFD session: **Up**, **Down**, **Init** (initializing), or **Failing**.	**brief detail extensive** none
Interface	Interface on which the BFD session is active.	**brief detail extensive** none
Detect Time	Negotiated time interval, in seconds, used to detect BFD control packets.	**brief detail extensive** none
Transmit Interval	Time interval, in seconds, used by the transmitting system to send BFD control packets.	**brief detail extensive** none
Multiplier	Negotiated multiplier by which the time interval is multiplied to determine the detection time for the transmitting system.	**detail extensive**
Session up time	How long a BFD session has been established.	**detail extensive**
Client	Protocol or process for which the BFD session is active: **ISIS**, **OSPF**, **DHCP**, **Static**, or **VGD**.	**detail extensive**
TX interval	Time interval, in seconds, used by the host system to transmit BFD control packets.	**brief detail extensive** none
RX interval	Time interval, in seconds, used by the host system to receive BFD control packets.	**brief detail extensive** none
Authenticate	Indicates that BFD authentication is configured.	**detail extensive**
keychain	Name of the security authentication keychain being used by a specific client. BFD authentication information for a client is provided in a single line and includes the **keychain**, **algo**, and **mode** parameters. Multiple clients can be configured on a BFD session.	**extensive**
algo	BFD authentication algorithm being used for a specific client: **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**. BFD authentication information for a client is provided in a single line and includes the **keychain**, **algo**, and **mode** parameters. Multiple clients can be configured on a BFD session.	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 87: show bfd session Output Fields (continued)

Field Name	Field Description	Level of Output
mode	Level of BFD authentication enforcement being used by a specific client: **strict** or **loose**. Strict enforcement indicates that authentication is configured at both ends of the session (the default). Loose enforcement indicates that one end of the session might not be authenticated. BFD authentication information for a client is provided in a single line and includes the **keychain**, **algo**, and **mode** parameters. Multiple clients can be configured on a BFD session.	**extensive**
Local diagnostic	Local diagnostic information about failing BFD sessions. Following are the expected values for Local Diagnostic output field: • **None**—No diagnostic • **CtlExpire**—Control detection time expired • **EchoExpire**—Echo detection time expired • **NbrSignal**—Neighbor signalled session down • **FwdPlaneReset**—Forwarding plane reset • **PathDown**—Path down • **ConcatPathDown**—Concatenated path down • **AdminDown**—Administratively down	**detail extensive**
Remote diagnostic	Remote diagnostic information about failing BFD sessions. Following are the expected values for Remote Diagnostic output field: • **None**—No diagnostic • **CtlExpire**—Control detection time expired • **EchoExpire**—Echo detection time expired • **NbrSignal**—Neighbor signalled session down • **FwdPlaneReset**—Forwarding plane reset • **PathDown**—Path down • **ConcatPathDown**—Concatenated path down • **AdminDown**—Administratively down	**detail extensive**
Remote state	Reports whether the remote system's BFD packets have been received and whether the remote system is receiving transmitted control packets.	**detail extensive**
Version	BFD version: **0** or **1**.	**extensive**
Replicated	The **replicated** flag appears when nonstop routing or graceful Routing Engine switchover is configured and the BFD session has been replicated to the backup Routing Engine.	**detail extensive**
Min async interval	Minimum amount of time, in seconds, between asynchronous control packet transmissions across the BFD session.	**extensive**
Min slow interval	Minimum amount of time, in seconds, between synchronous control packet transmissions across the BFD session.	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 87: show bfd session Output Fields (continued)

Field Name	Field Description	Level of Output
Adaptive async TX interval	Transmission interval being used because of adaptation.	extensive
RX interval	Minimum required receive interval.	extensive
Local min TX interval	Minimum amount of time, in seconds, between control packet transmissions on the local system.	extensive
Local min RX interval	Minimum amount of time, in seconds, between control packet detections on the local system.	extensive
Remote min TX interval	Minimum amount of time, in seconds, between control packet transmissions on the remote system.	extensive
Remote min TX interval	Minimum amount of time, in seconds, between control packet detections on the remote system.	extensive
Threshold transmission interval	Threshold for notification if the transmission interval increases.	extensive
Threshold for detection time	Threshold for notification if the detection time increases.	extensive
Local discriminator	Authentication code used by the local system to identify that BFD session.	extensive
Remote discriminator	Authentication code used by the remote system to identify that BFD session.	extensive
Echo mode	Information about the state of echo transmissions on the BFD session.	extensive
Prefix	LDP FEC address associated with the BFD session.	All levels
Egress, Destination	Displays the LDP FEC destination address. This field is displayed only on a router at the egress of an LDP FEC, where the BFD session has an LDP Operation, Administration, and Maintenance (OAM) client.	All levels
Remote is control-plane independent	The BFD session on the remote peer is running on its Packet Forwarding Engine. In this case, when the remote node undergoes a graceful restart, the local peer can help the remote peer with the graceful restart. The following BFD sessions are not distributed to the Packet Forwarding Engine: tunnel-encapsulated sessions, and sessions over integrated routing and bridging (IRB) interfaces.	extensive

 Copyright © 2018, Juniper Networks, Inc.

Table 87: show bfd session Output Fields (continued)

Field Name	Field Description	Level of Output
Authentication	Summary status of BFD authentication: • **status—enabled/active** indicates authentication is configured and active. **enabled/inactive** indicates authentication is configured but not active. This only occurs when the remote end of the session does not support authentication and loose checking is configured. • **keychain**—Name of the security authentication keychain associated with the specified BFD session. • **algo**—BFD authentication algorithm being used: **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**. • **mode**—Level of BFD authentication enforcement: **strict** or **loose**. Strict enforcement indicates authentication is configured at both ends of the session (the default). Loose enforcement indicates that one end of the session might not be authenticated. This information is only shown if BFD authentication is configured.	**extensive**
Session ID	The BFD session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA).	**detail extensive**
sessions	Total number of active BFD sessions.	All levels
clients	Total number of clients that are hosting active BFD sessions.	All levels
Cumulative transmit rate	Total number of BFD control packets transmitted per second on all active sessions.	All levels
Cumulative receive rate	Total number of BFD control packets received per second on all active sessions.	All levels
Multi-hop, min-recv-TTL	Minimum time to live (TTL) accepted if the session is configured for multihop.	**extensive**
route table	Route table used if the session is configured for multihop.	**extensive**
local address	Local address of the source used if the session is configured for multihop. The source IP address for outgoing BFD packets from the egress side of an MPLS BFD session is based on the outgoing interface IP address.	**extensive**

Sample Output

show bfd session

```
user@host> show bfd session
                                                      Transmit
Address              State     Interface     Detect Time  Interval  Multiplier
10.9.1.33            Up        so-7/1/0.0         0.600     0.200      3
10.9.1.29            Up        ge-4/0/0.0         0.600     0.200      3
```

```
2 sessions, 2 clients
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps
```

show bfd session brief

The output for the **show bfd session brief** command is identical to that for the **show bfd session** command.

show bfd session detail

```
user@host> show bfd session detail
                                                     Transmit
Address              State     Interface     Detect Time  Interval  Multiplier
10.9.1.33            Up        so-7/1/0.0          0.600     0.200      3
  Client OSPF, TX interval 0.200, RX interval 0.200, multiplier 3
  Session up time 3d 00:34:02
  Local diagnostic None, remote diagnostic None
  Remote state Up, version 1
  Replicated
10.9.1.29            Up        ge-4/0/0.0          0.600     0.200      3
  Client ISIS L2, TX interval 0.200, RX interval 0.200, multiplier 3
  Session up time 3d 00:29:04, previous down time 00:00:01
  Local diagnostic NbrSignal, remote diagnostic AdminDown
  Remote state Up, version 1

2 sessions, 2 clients
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps
```

show bfd session detail (with Authentication)

```
user@host>  show bfd session detail
                                                     Transmit
Address              State     Interface     Detect Time  Interval  Multiplier
10.9.1.33            Up        so-7/1/0.0          0.600     0.200      3
  Client OSPF, TX interval 0.200, RX interval 0.200, multiplier 3, Authenticate
  Session up time 3d 00:34:18
  Local diagnostic None, remote diagnostic None
  Remote state Up, version 1
  Replicated
10.9.1.29            Up        ge-4/0/0.0          0.600     0.200      3
  Client ISIS L2, TX interval 0.200, RX interval 0.200, multiplier 3
  Session up time 3d 00:29:12, previous down time 00:00:01
  Local diagnostic NbrSignal, remote diagnostic AdminDown
  Remote state Up, version 1

2 sessions, 2 clients
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps
```

show bfd session address extensive

```
user@host> show bfd session 10.255.245.212 extensive
                                                     Transmit
Address              State     Interface     Detect Time  Interval  Multiplier
10.255.245.212       Up                            1.200     0.400      3
  Client Static, TX interval 0.400, RX interval 0.400, multiplier 3
  Session up time 00:17:03, previous down time 00:00:14
  Local diagnostic CtlExpire, remote diagnostic NbrSignal
  Remote state Up, version 1
  Replicated
```


Copyright © 2018, Juniper Networks, Inc.

```
  Min async interval 0.400, min slow interval 1.000
  Adaptive async tx interval 0.400, rx interval 0.400
  Local min tx interval 0.400, min rx interval 0.400, multiplier 3
  Remote min tx interval 0.400, min rx interval 0.400, multiplier 3
  Threshold transmission interval 0.000, Threshold for detection time 0.000
  Local discriminator 6, remote discriminator 16
  Echo mode disabled/inactive
  Multi-hop, min-recv-TTL 255, route-table 0, local-address 10.255.245.205


1 sessions, 1 clients
Cumulative transmit rate 2.5 pps, cumulative receive rate 2.5 pps
```

show bfd session client rsvp-oam

```
user@host> show bfd session client rsvp-oam
                                             Detect   Transmit
 Address                  State     Interface      Time     Interval  Multiplier

 192.168.0.223            Up                       540.000   180.000      3

 1 Up sessions, 0 Down sessions
 1 sessions, 1 clients
 Cumulative transmit rate 0.0 pps, cumulative receive rate 0.0 pps
```

show bfd session client vpls-oam summary

```
user@host> show bfd session client vpls-oam summary
1 Up sessions, 1 Down sessions
 2 sessions, 2 clients
 Cumulative transmit rate 2.0 pps, cumulative receive rate 1.0 pps
```

show bfd session client vpls-oam instance instance-name

```
user@host> show bfd session client vpls-oam instance vpls
                                             Detect   Transmit
 Address                  State     Interface      Time     Interval  Multiplier

 127.0.0.1                Up        ae9.0          3.000     1.000        3

 1 Up Sessions, 0 Down Sessions
 1 sessions, 1 clients
 Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps
```

show bfd session extensive

```
user@host> show bfd session extensive
                                            Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
10.31.1.2                Up        ge-2/1/8.0     0.030     0.010        3
 Client OSPF realm ospf-v2 Area 0.0.0.0, TX interval 0.010, RX interval 0.010
 Session up time 00:10:13
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Replicated
 Min async interval 0.010, min slow interval 1.000
 Adaptive async TX interval 0.010, RX interval 0.010
 Local min TX interval 0.010, minimum RX interval 0.010, multiplier 3
 Remote min TX interval 0.010, min RX interval 0.010, multiplier 3
 Local discriminator 12, remote discriminator 4
```

Copyright © 2018, Juniper Networks, Inc.

```
 Echo mode disabled/inactive
 Remote is control-plane independent
  Session ID: 0x201
  Micro-BFD Session
                                                 Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
10.31.2.2                Up        ge-2/1/4.0     0.030     0.010        3
 Client OSPF realm ospf-v2 Area 0.0.0.0, TX interval 0.010, RX interval 0.010
 Session up time 00:10:14
 Local diagnostic None, remote diagnostic NbrSignal
 Remote state Up, version 1
 Replicated
 Min async interval 0.010, min slow interval 1.000
 Adaptive async TX interval 0.010, RX interval 0.010
 Local min TX interval 0.010, minimum RX interval 0.010, multiplier 3
 Remote min TX interval 0.010, min RX interval 0.010, multiplier 3
 Local discriminator 13, remote discriminator 5
 Echo mode disabled/inactive
 Remote is control-plane independent
  Session ID: 0x202

2 sessions, 2 clients
Cumulative transmit rate 200.0 pps, cumulative receive rate 200.0 pps
```

show bfd session extensive (with Authentication)

```
user@host> show bfd session extensive
                                            Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
192.168.208.26           Up        so-1/0/0.0     2.400     0.800        10
 Client Static, TX interval 0.600, RX interval 0.600, Authenticate
       keychain bfd, algo keyed-md5, mode loose
 Session up time 00:18:07
 Local diagnostic None, remote diagnostic NbrSignal
 Remote state Up, version 1
 Replicated
 Min async interval 0.600, min slow interval 1.000
 Adaptive async TX interval 0.600, RX interval 0.600
 Local min TX interval 0.600, minimum RX interval 0.600, multiplier 10
 Remote min TX interval 0.800, min RX interval 0.800, multiplier 3
 Local discriminator 2, remote discriminator 3
 Echo mode disabled/inactive
 Authentication enabled/active, keychain bfd, algo keyed-md5, mode loose


1 sessions, 1 clients
Cumulative transmit rate 1.2 pps, cumulative receive rate 1.2 pps
```

show bfd session summary

```
user@host> show bfd session summary
2 sessions, 2 clients
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps
```

show bfd session subscriber

```
user@host> show bfd session subscriber
      Detect  Transmit
Address   State  Interface  Time  Interval  Multiplier
1.0.0.2   Up   ae0.0  90.000  30.000   3
```

 Copyright © 2018, Juniper Networks, Inc.

```
1.0.0.6   Up  ae0.1  90.000  30.000  3
1.0.0.10 Up  ae0.2  90.000 30.000   3
1.0.0.14 Up  ae0.3  90.000  30.000   3
1.0.0.18 Up   ae0.4  90.000  30.000   3

20 sessions, 20 clients
```

show bfd session subscriber address

```
user@host> show bfd session subscriber address 1.0.0.2
     Detect  Transmit
Address  State  Interface Time  Interval  Multiplier
1.0.0.2  Up   ae0.0  90.000  30.000   3

1 sessions, 1 clients
Cumulative transmit rate 5.0 pps, cumulative receive rate 5.0 pps
```

show bfd session subscriber extensive

```
user@host> show bfd session subscriber extensive
     Detect   Transmit
Address  State  Interface Time   Interval  Multiplier
1.0.0.2  Up   ae0.0  90.000   30.000   3

Client DHCP, TX interval 30.000, RX interval 30.000
Session up time 09:11:50
Local diagnostic None, remote diagnostic NbrSignal
Remote state Up, version 1
Replicated
Min async interval 30.000, min slow interval 30.000
Adaptive async TX interval 30.000, RX interval 30.000
Local min TX interval 30.000, minimum RX interval 30.000, multiplier 3
Remote min TX interval 30.000, min RX interval 30.000, multiplier 3
Local discriminator 20, remote discriminator 16
Echo mode disabled/inactive
Remote is control-plane independent
Session ID: 0x1

     Detect   Transmit
Address  State  Interface  Time   Interval  Multiplier
1.0.0.6  Up   ae0.1  90.000   30.000   3

Client DHCP, TX interval 30.000, RX interval 30.000
Session up time 09:11:50
Local diagnostic None, remote diagnostic NbrSignal
Remote state Up, version 1
Replicated
Min async interval 30.000, min slow interval 30.000
Adaptive async TX interval 30.000, RX interval 30.000
Local min TX interval 30.000, minimum RX interval 30.000, multiplier 3
Remote min TX interval 30.000, min RX interval 30.000, multiplier 3
Local discriminator 21, remote discriminator 17
Echo mode disabled/inactive
Remote is control-plane independent
Session ID: 0x2
```

show bfd session subscriber discriminator extensive

```
user@host> show bfd session subscriber discriminator 20 extensive
```

Copyright © 2018, Juniper Networks, Inc.

```
     Detect   Transmit
Address  State  Interface  Time   Interval  Multiplier
1.0.0.2  Up   ae0.0  90.000   30.000   3

Client DHCP, TX interval 30.000, RX interval 30.000
Session up time 09:11:50
Local diagnostic None, remote diagnostic NbrSignal
Remote state Up, version 1
Replicated
Min async interval 30.000, min slow interval 30.000
Adaptive async TX interval 30.000, RX interval 30.000
Local min TX interval 30.000, minimum RX interval 30.000, multiplier 3
Remote min TX interval 30.000, min RX interval 30.000, multiplier 3
Local discriminator 20, remote discriminator 16
Echo mode disabled/inactive
Remote is control-plane independent
Session ID: 0x1

1 sessions, 1 clients
Cumulative transmit rate 5.0 pps, cumulative receive rate 5.0 pps
```

 Copyright © 2018, Juniper Networks, Inc.

CHAPTER 17

BGP Operational Commands

clear bgp damping

List of Syntax

Syntax on page 1664
Syntax (EX Series Switch and QFX Series) on page 1664

Syntax

clear bgp damping
<logical-system (all | *logical-system-name*)>
<*prefix*>

Syntax (EX Series Switch and QFX Series)

clear bgp damping
<*prefix*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Clear BGP route flap damping information.

Options

none—Clear all BGP route flap damping information.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

prefix—(Optional) Clear route flap damping information for only the specified destination prefix.

Required Privilege Level

clear

Related Documentation

- show policy damping on page 1710
- show route damping on page 2215

List of Sample Output

clear bgp damping on page 1664

Output Fields

This command produces no output.

Sample Output

clear bgp damping

```
user@host> clear bgp damping
```


Copyright © 2018, Juniper Networks, Inc.

clear bgp neighbor

List of Syntax

Syntax on page 1665
Syntax (EX Series Switch and QFX Series) on page 1665

Syntax

clear bgp neighbor
<all>
<as *as-number*>
<gracefully>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<malformed-route>
<*neighbor*>
<soft | soft-inbound>
<soft-minimum-igp>
<stale-routes>

Syntax (EX Series Switch and QFX Series)

clear bgp neighbor
<all>
<as *as-number*>
<instance *instance-name*>
<malformed-route>
<*neighbor*>
<soft | soft-inbound>
<soft-minimum-igp>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
malformed-route option introduced in Junos OS Release 13.2.
all option introduced in Junos OS Release 14.2.
gracefully and **stale-routes** options introduced in Junos OS Release 15.1.

Description

Perform one of the following tasks:

- Change the state of one or more BGP neighbors to **IDLE**. For neighbors in the **ESTABLISHED** state, this command drops the TCP connection to the neighbors and then reestablishes the connection.
- (**soft** keyword only) Reapply export policies and send refresh updates to one or more BGP neighbors without changing their state.
- (**soft-inbound** keyword only) Send a route-refresh message to one or more BGP neighbors without changing their state, and reapply import policies on the received updates.

Options

all—Change the state of all BGP neighbors to **IDLE**.

as ***as-number***—(Optional) Apply this command only to neighbors in the specified autonomous system (AS).

gracefully—(Optional) Enable the BGP peer to start graceful-restart receiving-speaker mode. The receiving speaker also sends its own routes to the restarted speaker, and sends an End-of-RIB marker when it completes the update. The **clear bgp neighbor** ***neighbor-address*** **gracefully** command is the same as **clear bgp neighbor hard** (the default for **clear bgp neighbor**), but it does not use the new Hard Reset subcode on the Notify and Cease messages that are sent. This allows the neighbor to enter GR or LLGR helper mode, if negotiated. The session is still cleared on this router, and this router does not enter GR or LLGR helper mode.

instance ***instance-name***—(Optional) Apply this command only to neighbors for the specified routing instance.

logical-system (all | ***logical-system-name*****)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

malformed-route—(Optional) Remove malformed routes. If a specific neighbor is provided, Junos OS removes malformed routes for that particular neighbor. Otherwise, Junos OS removes malformed routes for all BGP neighbors. To find routes that have malformed attributes, run the **show route hidden** command, and look for routes marked with **MalformedAttr** in the AS path field.

neighbor—(Optional) IP address of a BGP peer. Apply this command only to the specified neighbor.

soft—(Optional) Reapply any export policies and send refresh updates to neighbors without clearing the state.

soft-inbound—(Optional) Send a route-refresh message to BGP neighbors and reapply import policies on the route updates received from the BGP neighbors without clearing the BGP state.

soft-minimum-igp—(Optional) Provide soft refresh of the outbound state when the interior gateway protocol (IGP) metric is reset.

stale-routes—(Optional) Any stale route currently being held for the specified neighbor because of BGP graceful restart (GR) or long-lived graceful restart (LLGR) receiver mode operations.

Required Privilege Level clear

Related Documentation

- show bgp neighbor on page 1683

List of Sample Output clear bgp neighbor on page 1667

Output Fields When you enter this command, you are provided feedback on the status of your request.

 Copyright © 2018, Juniper Networks, Inc.

Sample Output

clear bgp neighbor

`user@host>` **clear bgp neighbor**

clear bgp table

Syntax

clear bgp table ***table-name***
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and QFX Series)

clear bgp table ***table-name***

Release Information

Command introduced in Junos OS Release 9.0.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Request that BGP refresh routes in a specified routing table.

Options

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

table-name—Request that BGP refresh routes in the specified table.

Additional Information

In some cases, a prefix limit is associated with a routing table for a VPN instance. When this limit is exceeded (for example, because of a network misconfiguration), some routes might not be inserted in the table. Such routes need to be added to the table after the network issue is resolved. Use the **clear bgp table** command to request that BGP refresh routes in a VPN instance table.

Required Privilege Level

clear

List of Sample Output

clear bgp table private.inet.0 on page 1668
clear bgp table inet.6 logical-system all on page 1668
clear bgp table private.inet.6 logical-system ls1 on page 1669
clear bgp table logical-system all inet.0 on page 1669
clear bgp table logical-system ls2 private.inet.0 on page 1669

Output Fields

This command produces no output.

Sample Output

clear bgp table private.inet.0

```
user@host> clear bgp table private.inet.0
```

clear bgp table inet.6 logical-system all

```
user@host> clear bgp table inet.6 logical-system all
```


Copyright © 2018, Juniper Networks, Inc.

clear bgp table private.inet.6 logical-system ls1

`user@host>` **clear bgp table private.inet.6 logical-system ls1**

clear bgp table logical-system all inet.0

`user@host>` **clear bgp table logical-system all inet.0**

clear bgp table logical-system ls2 private.inet.0

`user@host>` **clear bgp table logical-system ls2 private.inet.0**

show bgp bmp

Syntax

show bgp bmp

Release Information

Command introduced in Junos OS Release 9.5.
Command introduced in Junos OS Release 9.5 for EX Series switches.
Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display information about the BGP Monitoring Protocol (BMP).

Options

This command has no options.

Required Privilege Level

view

List of Sample Output

show bgp bmp on page 1670

Output Fields

Table 88 on page 1670 lists the output fields for the **show bgp bmp** command. Output fields are listed in the approximate order in which they appear.

Table 88: show bgp bmp Output Fields

Field Name	Field Description
BMP station address/port	IP address and port number of the monitoring station to which BGP Monitoring Protocol (BMP) statistics are sent.
BMP session state	Status of the BMP session: **UP** or **DOWN**.
Memory consumed by BMP	Memory used by the active BMP session.
Statistics timeout	Amount of time, in seconds, between transmissions of BMP data to the monitoring station.
Memory limit	Threshold, in bytes, at which the routing device stops collecting BMP data.
Memory-connect retry timeout	Amount of time, in seconds, after which the routing device attempts to resume a BMP session that was ended after the configured memory threshold was exceeded.

Sample Output

show bgp bmp

```
user@host> show bgp bmp
  BMP station address/port: 172.24.24.157+5454
  BMP session state: DOWN
```


Copyright © 2018, Juniper Networks, Inc.

```
Memory consumed by BMP: 0
Statistics timeout: 15
Memory limit: 10485760
Memory connect retry timeout: 600
```

show bgp group

List of Syntax

Syntax on page 1672
Syntax (EX Series Switch and QFX Series) on page 1672

Syntax

show bgp group
<brief | detail | summary>
<*group-name*>
<exact-instance *instance-name*>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<rtf>

Syntax (EX Series Switch and QFX Series)

show bgp group
<brief | detail | summary>
<*group-name*>
<exact-instance *instance-name*>
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
exact-instance option introduced in Junos OS Release 11.4.

Description

Display information about the configured BGP groups.

Options

none—Display group information about all BGP groups.

brief | detail | summary—(Optional) Display the specified level of output.

group-name—(Optional) Display group information for the specified group.

exact-instance *instance-name*—(Optional) Display information for the specified instance only.

instance *instance-name*—(Optional) Display information about BGP groups for all routing instances whose name begins with this string (for example, **cust1**, **cust11**, and **cust111** are all displayed when you run the **show bgp group instance cust1** command). The instance name can be master for the main instance, or any valid configured instance name or its prefix.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

rtf—(Optional) Display BGP group route targeting information.

Required Privilege Level

view

 Copyright © 2018, Juniper Networks, Inc.

List of Sample Output

show bgp group on page 1676
show bgp group on page 1676
show bgp group brief on page 1677
show bgp group detail on page 1677
show bgp group rtf detail on page 1678
show bgp group summary on page 1678

Output Fields

Table 89 on page 1673 describes the output fields for the **show bgp group** command. Output fields are listed in the approximate order in which they appear.

Table 89: show bgp group Output Fields

Field Name	Field Description	Level of Output
Group Type or **Group**	Type of BGP group: **Internal** or **External**.	All levels
group-index	Index number for the BGP peer group. The index number differentiates between groups when a single BGP group is split because of different configuration options at the group and peer levels.	**rtf detail**
AS	AS number of the peer. For internal BGP (IBGP), this number is the same as **Local AS**.	**brief detail** none
Local AS	AS number of the local routing device.	**brief detail** none
Name	Name of a specific BGP group.	**brief detail** none
Options	The Network Layer Reachability Information (NLRI) format used for BGP VPN multicast.	**none** none
Index	Unique index number of a BGP group.	**brief detail** none
Flags	Flags associated with the BGP group. This field is used by Juniper Networks customer support.	**brief detail** none
BGP-Static Advertisement Policy	Policies configured for the BGP group with the **advertise-bgp-static policy** statement.	**brief** none
Remove-private options	Options associated with the **remove-private** statement.	**brief detail** none
Holdtime	Maximum number of seconds allowed to elapse between successive keepalive or update messages that BGP receives from a peer in the BGP group, after which the connection to the peer is closed and routing devices through that peer become unavailable.	**brief detail** none
Export	Export policies configured for the BGP group with the **export** statement.	**brief detail** none

Copyright © 2018, Juniper Networks, Inc.

Table 89: show bgp group Output Fields (continued)

Field Name	Field Description	Level of Output
Optimal Route Reflection	Client nodes (primary and backup) configured in the BGP group.	**brief detail** none
MED tracks IGP metric update delay	Time, in seconds, that updates to multiple exit discriminator (MED) are delayed. Also displays the time remaining before the interval is set to expire	All levels
Traffic Statistics Interval	Time between sample periods for labeled-unicast traffic statistics, in seconds.	**brief detail** none
Total peers	Total number of peers in the group.	**brief detail** none
Established	Number of peers in the group that are in the established state.	All levels

Active/Received/Accepted/Damped	Multipurpose field that displays information about BGP peer sessions. The field's contents depend upon whether a session is established and whether it was established in the main routing device or in a routing instance.
	• If a peer is not established, the field shows the state of the peer session: **Active**, **Connect**, or **Idle**.
	• If a BGP session is established in the main routing device, the field shows the number of active, received, accepted, and damped routes that are received from a neighbor and appear in the **inet.0** (main) and **inet.2** (multicast) routing tables. For example, **8/10/10/2 and 2/4/4/0** indicate the following:
	  ◦ 8 active routes, 10 received routes, 10 accepted routes, and 2 damped routes from a BGP peer appear in the **inet.0** routing table.
	  ◦ 2 active routes, 4 received routes, 4 accepted routes, and no damped routes from a BGP peer appear in the **inet.2** routing table.
		summary

ip-addresses	List of peers who are members of the group. The address is followed by the peer's port number.	All levels
Route Queue Timer	Number of seconds until queued routes are sent. If this time has already elapsed, this field displays the number of seconds by which the updates are delayed.	**detail**
Route Queue	Number of prefixes that are queued up for sending to the peers in the group.	**detail**
inet.*number*	Number of active, received, accepted, and damped routes in the routing table. For example, **inet.0: 7/10/9/0** indicates the following: • 7 active routes, 10 received routes, 9 accepted routes, and no damped routes from a BGP peer appear in the **inet.0** routing table.	none

 Copyright © 2018, Juniper Networks, Inc.

Table 89: show bgp group Output Fields (continued)

Field Name	Field Description	Level of Output
Table inet.*number*	Information about the routing table. • **Received prefixes**—Total number of prefixes from the peer, both active and inactive, that are in the routing table. • **Active prefixes**—Number of prefixes received from the peer that are active in the routing table. • **Suppressed due to damping**—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols. • **Advertised prefixes**—Number of prefixes advertised to a peer. • **Received external prefixes**—Total number of prefixes from the external BGP (EBGP) peers, both active and inactive, that are in the routing table. • **Active external prefixes**—Number of prefixes received from the EBGP peers that are active in the routing table. • **Externals suppressed**—Number of routes received from EBGP peers currently inactive because of damping or other reasons. • **Received internal prefixes**—Total number of prefixes from the IBGP peers, both active and inactive, that are in the routing table. • **Active internal prefixes**—Number of prefixes received from the IBGP peers that are active in the routing table. • **Internals suppressed**—Number of routes received from IBGP peers currently inactive because of damping or other reasons. • **RIB State**—Status of the graceful restart process for this routing table: **BGP restart is complete**, **BGP restart in progress**, **VPN restart in progress**, or **VPN restart is complete**.	**detail**
Groups	Total number of groups.	All levels
Peers	Total number of peers.	All levels
External	Total number of external peers.	All levels
Internal	Total number of internal peers.	All levels
Down peers	Total number of unavailable peers.	All levels
Flaps	Total number of flaps that occurred.	All levels
Table	Name of a routing table.	**brief**, none
Tot Paths	Total number of routes.	**brief**, none
Act Paths	Number of active routes.	**brief**, none
Suppressed	Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.	**brief**, none

Copyright © 2018, Juniper Networks, Inc.

Table 89: show bgp group Output Fields (continued)

Field Name	Field Description	Level of Output
History	Number of withdrawn routes stored locally to keep track of damping history.	**brief**, none
Damp State	Number of active routes with a figure of merit greater than zero, but lower than the threshold at which suppression occurs.	**brief**, none
Pending	Routes being processed by the BGP import policy.	**brief**, none
Group	Group the peer belongs to in the BGP configuration.	**detail**
Receive mask	Mask of the received target included in the advertised route.	**detail**
Entries	Number of route entries received.	**detail**
Target	Route target that is to be passed by route-target filtering. If a route advertised from the provider edge (PE) routing device matches an entry in the route-target filter, the route is passed to the peer.	**detail**
Mask	Mask which specifies that the peer receive routes with the given route target.	**detail**

Sample Output

show bgp group

```
user@host> show bgp group
```

show bgp group

```
user@host> show bgp group
Group Type: Internal    AS: 1001                   Local AS: 1001
  Name: ibgp            Index: 2                   Flags: Export Eval
  Holdtime: 0
  Optimal Route Reflection: igp-primary 1.1.1.1, igp-backup 1.1.2.1
  Total peers: 1        Established: 1
  1.1.1.2+179
  Trace options:  all
  Trace file: /var/log/bgp-log size 10485760 files 10
  bgp.l3vpn.2: 0/0/0/0
  vpn-1.inet.2: 0/0/0/0

Group Type: Internal    AS: 1001                   Local AS: 1001
  Name: ibgp            Index: 3                   Flags: Export Eval
  Options: RFC6514CompliantSafi129
  Holdtime: 0
  Optimal Route Reflection: igp-primary 1.1.1.1, igp-backup 1.1.2.1
  Total peers: 1        Established: 1
  1.1.1.5+61698
  Trace options:  all
  Trace file: /var/log/bgp-log size 10485760 files 10
  bgp.l3vpn.2: 2/2/2/0
```

 Copyright © 2018, Juniper Networks, Inc.

```
  vpn-1.inet.2: 2/2/2/0

Groups: 2  Peers: 2    External: 0    Internal: 2    Down peers: 0   Flaps: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
bgp.l3vpn.2
                       2          2          0          0          0          0
vpn-1.inet.0
                       0          0          0          0          0          0
vpn-1.inet.2
                       2          2          0          0          0          0
vpn-1.inet6.0
                       0          0          0          0          0          0
vpn-1.mdt.0
                       0          0          0          0          0          0
```

show bgp group brief

```
user@host> show bgp group brief
 Groups: 2  Peers: 2    External: 0    Internal: 2    Down peers: 1   Flaps: 0
 Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending

 inet.0
                       0          0          0          0          0          0

 bgp.l3vpn.0
                       0          0          0          0          0          0

 bgp.rtarget.0
                       2          0          0          0          0          0
```

show bgp group detail

```
user@host> show bgp group detail
Group Type: Internal    AS: 1                      Local AS: 1
  Name: ibgp            Index: 0                   Flags: <Export Eval>
  Holdtime: 0
  Optimal Route Reflection: igp-primary 1.1.1.1, igp-backup 1.1.2.1
  Total peers: 3        Established: 0
  22.0.0.2
  22.0.0.8
  22.0.0.5

Groups: 1  Peers: 3    External: 0    Internal: 3    Down peers: 3   Flaps: 3
  Table bgp.l3vpn.0
    Received prefixes:            0
    Accepted prefixes:            0
    Active prefixes:              0
    Suppressed due to damping:    0
    Received external prefixes:   0
    Active external prefixes:     0
    Externals suppressed:         0
    Received internal prefixes:   0
    Active internal prefixes:     0
    Internals suppressed:         0
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
  Table bgp.mdt.0
    Received prefixes:            0
    Accepted prefixes:            0
    Active prefixes:              0
```

Copyright © 2018, Juniper Networks, Inc.

```
    Suppressed due to damping:    0
    Received external prefixes:   0
    Active external prefixes:     0
    Externals suppressed:         0
    Received internal prefixes:   0
    Active internal prefixes:     0
    Internals suppressed:         0
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
  Table VPN-A.inet.0
    Received prefixes:            0
    Accepted prefixes:            0
    Active prefixes:              0
    Suppressed due to damping:    0
    Received external prefixes:   0
    Active external prefixes:     0
    Externals suppressed:         0
    Received internal prefixes:   0
    Active internal prefixes:     0
    Internals suppressed:         0
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
  Table VPN-A.mdt.0
    Received prefixes:            0
    Accepted prefixes:            0
    Active prefixes:              0
    Suppressed due to damping:    0
    Received external prefixes:   0
    Active external prefixes:     0
    Externals suppressed:         0
    Received internal prefixes:   0
    Active internal prefixes:     0
    Internals suppressed:         0
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
```

show bgp group rtf detail

```
user@host> show bgp group rtf detail
 Group: internal (group-index: 0)
     Receive mask: 00000002
     Table: bgp.rtarget.0                                   Entries: 2
         Target                          Mask
         100:100/64                      00000002
         200:201/64                      (Group)
 Group: internal (group-index: 1)
     Table: bgp.rtarget.0                                   Entries: 1
         Target                          Mask
         200:201/64                      (Group)
```

show bgp group summary

```
user@host> show bgp group summary
Group        Type       Peers     Established    Active/Received/Accepted/Damped
ibgp         Internal   3         0

Groups: 1  Peers: 3    External: 0    Internal: 3    Down peers: 3    Flaps: 3
  bgp.l3vpn.0       : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
  bgp.mdt.0         : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
```

 Copyright © 2018, Juniper Networks, Inc.

```
VPN-A.inet.0     : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
VPN-A.mdt.0      : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
```

Copyright © 2018, Juniper Networks, Inc.

show bgp group traffic-statistics

Syntax

show bgp group traffic-statistics
<brief | detail>
<*group-name*>
<labeled-path label *label*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced before Junos OS Release 7.4.
labeled-path option introduced in Junos OS Release 18.1R1 for the MX Series.

Description

Display the traffic statistics for configured Border Gateway Protocol (BGP) groups.

Options

none—Display traffic statistics for all BGP groups.

brief | detail—(Optional) Display the specified level of output.

group-name—(Optional) Display BGP traffic statistics for only the specified group.

label-path—(Optional) Display labeled unicast traffic statistics at the ingress.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show bgp group traffic-statistics (Per-Group-Label Not Configured) on page 1681
show bgp group traffic-statistics (Per-Group-Label Configured) on page 1681
show bgp group traffic-statistics labeled-path (Labeled Unicast) on page 1682

Output Fields

Table 90 on page 1680 describes the output fields for the **show bgp group traffic-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 90: show bgp group traffic-statistics Output Fields

Field Name	Field Description
Group name	Name of a specific BGP group.
Group Index	Index number for the BGP group.
NLRI	Network layer reachability information (NLRI) indicating the source of the traffic statistics for the BGP group.
FEC	Forwarding equivalence classes (FECs) associated with the BGP group.
Packets	Number of packets sent through each FEC.
Bytes	Number of bytes transmitted through each FEC.

Copyright © 2018, Juniper Networks, Inc.

Table 90: show bgp group traffic-statistics Output Fields (continued)

Field Name	Field Description
EgressAS	Autonomous system (AS) number of the egress router.
AdvLabel	Label associated with each FEC.

Sample Output

show bgp group traffic-statistics (Per-Group-Label Not Configured)

```
user@host> show bgp group traffic-statistics
Group Name: ext1        Group Index: 0          NLRI: inet-labeled-unicast
FEC                    Packets                Bytes     EgressAS     AdvLabel
10.255.245.55                0                    0            I       100224
10.255.245.57                0                    0            I       100240
100.101.0.0                550                48400           25       100256
100.102.0.0                550                48400           25       100256
100.103.0.0                550                48400           25       100272
100.104.0.0                550                48400           25       100272
192.168.25.0                 0                    0            I       100288

Group Name: ext2        Group Index: 1          NLRI: inet-labeled-unicast
FEC                    Packets                Bytes     EgressAS     AdvLabel
10.255.245.55                0                    0            I       100224
10.255.245.57                0                    0            I       100240
100.101.0.0                550                48400           25       100256
100.102.0.0                550                48400           25       100256
100.103.0.0                550                48400           25       100272
100.104.0.0                550                48400           25       100272
192.168.25.0                 0                    0            I       100288
```

show bgp group traffic-statistics (Per-Group-Label Configured)

```
user@host> show bgp group traffic-statistics
Group Name: ext1        Group Index: 0          NLRI: inet-labeled-unicast
FEC                    Packets                Bytes     EgressAS     AdvLabel
10.255.245.55                0                    0            I       100384
10.255.245.57                0                    0            I       100400
100.101.0.0                101                 8888           25       100416
100.102.0.0                101                 8888           25       100416
100.103.0.0                  0                    0           25       100432
100.104.0.0                  0                    0           25       100432
192.168.25.0                 0                    0            I       100448

Group Name: ext2        Group Index: 1          NLRI: inet-labeled-unicast
FEC                    Packets                Bytes     EgressAS     AdvLabel
10.255.245.55                0                    0            I       100304
10.255.245.57                0                    0            I       100320
100.101.0.0                  0                    0           25       100336
100.102.0.0                  0                    0           25       100336
100.103.0.0                101                 8888           25       100352
100.104.0.0                101                 8888           25       100352
192.168.25.0                 0                    0            I       100368
```

show bgp group traffic-statistics labeled-path (Labeled Unicast)

```
user@host> show bgp group traffic-statistics labeled-path
Labels                  NextHop             Packets               Bytes
3(top)                 10.1.1.1                   0                   0
299840(top)            40.1.1.1                   0                   0
110001(top)            40.1.1.1                   2                 168
110002
110003
110001(top)            40.1.1.1                   0                   0
110072
110073
110071(top)            40.1.1.1                   0                   0
110072
110073
120001(top)            40.1.1.1                   0                   0
120002
120003
1000002(top)           40.1.1.1                   2                 168
1000003
1000004
```

 Copyright © 2018, Juniper Networks, Inc.

show bgp neighbor

List of Syntax

Syntax on page 1683
Syntax (EX Series Switch, QFX Series, and OCX Series) on page 1683

Syntax

show bgp neighbor
<exact-instance *instance-name*>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<*neighbor-address*>
<output-queue>
<orf (detail | *neighbor-address*)

Syntax (EX Series Switch, QFX Series, and OCX Series)

show bgp neighbor
<instance *instance-name*>
<exact-instance *instance-name*>
<*neighbor-address*>
<orf (*neighbor-address* | detail)

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1x53-D20 for the OCX Series.
orf option introduced in Junos OS Release 9.2.
exact-instance option introduced in Junos OS Release 11.4.
output-queue option introduced in Junos OS Release 16.1

Description

Display information about BGP peers.

Options

none—Display information about all BGP peers.

exact-instance ***instance-name***—(Optional) Display information for the specified instance only.

instance ***instance-name***—(Optional) Display information about BGP peers for all routing instances whose name begins with this string (for example, **cust1**, **cust11**, and **cust111** are all displayed when you run the **show bgp neighbor instance cust1** command).

logical-system (all | ***logical-system-name*****)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

neighbor-address—(Optional) Display information for only the BGP peer at the specified IP address.

orf (detail | ***neighbor-address*****)**—(Optional) Display outbound route-filtering information for all BGP peers or only for the BGP peer at the specified IP address. The default is to display brief output. Use the **detail** option to display detailed output.

output-queue—(Optional) Display information regarding the number of routes currently queued in the 17 prioritized BGP output queues.

Copyright © 2018, Juniper Networks, Inc.

Additional Information For information about the **local-address, nlri**, **hold-time,** and **preference** statements, see the *Junos OS Routing Protocols Library*.

Required Privilege Level view

Related Documentation

- clear bgp neighbor on page 1665

List of Sample Output

Output Fields Table 91 on page 1684 describes the output fields for the **show bgp neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 91: show bgp neighbor Output Fields

Field Name	Field Description
Peer	Address of the BGP neighbor. The address is followed by the neighbor port number.
AS	AS number of the peer.
Local	Address of the local routing device. The address is followed by the peer port number.
Type	Type of peer: **Internal** or **External**.
State	Current state of the BGP session: • **Active**—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message. • **Connect**—BGP is waiting for the transport protocol connection to be completed. • **Established**—The BGP session has been established, and the peers are exchanging update messages. • **Idle**—This is the first stage of a connection. BGP is waiting for a Start event. • **OpenConfirm**—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • **OpenSent**—BGP has sent an open message and is waiting to receive an open message from the peer. • **route reflector client**—The BGP session is established with a route reflector client.

 Copyright © 2018, Juniper Networks, Inc.

Table 91: show bgp neighbor Output Fields (continued)

Field Name	Field Description
Flags	Internal BGP flags: • **Aggregate Label**—BGP has aggregated a set of incoming labels (labels received from the peer) into a single forwarding label. • **CleanUp**—The peer session is being shut down. • **Delete**—This peer has been deleted. • **Idled**—This peer has been permanently idled. • **ImportEval**—At the last commit operation, this peer was identified as needing to reevaluate all received routes. • **Initializing**—The peer session is initializing. • **SendRtn**—Messages are being sent to the peer. • **Sync**—This peer is synchronized with the rest of the peer group. • **RSync**—This peer in the backup Routing Engine is synchronized with the BGP peer in the master Routing Engine for nonstop active routing. • **TryConnect**—Another attempt is being made to connect to the peer. • **Unconfigured**—This peer is not configured. • **WriteFailed**—An attempt to write to this peer failed.
Last state	Previous state of the BGP session: • **Active**—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message. • **Connect**—BGP is waiting for the transport protocol connection to be completed. • **Established**—The BGP session has been established, and the peers are exchanging update messages. • **Idle**—This is the first stage of a connection. BGP is waiting for a Start event. • **OpenConfirm**—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • **OpenSent**—BGP has sent an open message and is waiting to receive an open message from the peer.
Last event	Last activity that occurred in the BGP session: • **Closed**—The BGP session closed. • **ConnectRetry**—The transport protocol connection failed, and BGP is trying again to connect. • **HoldTime**—The session ended because the hold timer expired. • **KeepAlive**—The local routing device sent a BGP keepalive message to the peer. • **Open**—The local routing device sent a BGP open message to the peer. • **OpenFail**—The local routing device did not receive an acknowledgment of a BGP open message from the peer. • **RecvKeepAlive**—The local routing device received a BGP keepalive message from the peer. • **RecvNotify**—The local routing device received a BGP notification message from the peer. • **RecvOpen**—The local routing device received a BGP open message from the peer. • **RecvUpdate**—The local routing device received a BGP update message from the peer. • **Start**—The peering session started. • **Stop**—The peering session stopped. • **TransportError**—A TCP error occurred.

Copyright © 2018, Juniper Networks, Inc.

Table 91: show bgp neighbor Output Fields (continued)

Field Name	Field Description
Last error	Last error that occurred in the BGP session: • **Cease**—An error occurred, such as a version mismatch, that caused the session to close. • **Finite State Machine Error**—In setting up the session, BGP received a message that it did not understand. • **Hold Time Expired**—The session's hold time expired. • **Message Header Error**—The header of a BGP message was malformed. • **Open Message Error**—A BGP open message contained an error. • **None**—No errors occurred in the BGP session. • **Update Message Error**—A BGP update message contained an error.
Export	Name of the export policy that is configured on the peer.
Import	Name of the import policy that is configured on the peer.
Options	Configured BGP options: • **AddressFamily**—Configured address family: inet or inet-vpn. • **AdvertiseBGPStatic**—Configured BGP static routes are advertised. • **AutheKeyChain**—Authentication key change is enabled. • **DropPathAttributes**—Certain path attributes are configured to be dropped from neighbor updates during inbound processing. • **GracefulRestart**—Graceful restart is configured. • **HoldTime**—Hold time configured with the **hold-time** statement. The hold time is three times the interval at which keepalive messages are sent. • **IgnorePathAttributes**—Certain path attributes are configured to be ignored in neighbor updates during inbound processing. • **Local Address**—Address configured with the **local-address** statement. • **LLGR**—BGP long-lived graceful restart capability is configured. • **LLGRHelperDisabled**—BGP long-lived graceful restart is completely disabled for a neighbor. • **Multihop**—Allow BGP connections to external peers that are not on a directly connected network. • **NLRI**—Configured MBGP state for the BGP group: **multicast**, **unicast**, or both if you have configured **nlri any**. • **Peer AS**—Configured peer autonomous system (AS). • **Preference**—Preference value configured with the **preference** statement. • **Refresh**—Configured to refresh automatically when the policy changes. • **Rib-group**—Configured routing table group. • **RFC6514CompliantSafi129**—Configured SAFI 129 according to RFC 6514 (BGP VPN multicast used to use SAFI 128).
Path-attributes dropped	Path attribute codes that are dropped from neighbor updates.
Path-attributes ignored	Path attribute codes that are ignored during neighbor updates.
Peer does not support LLGR Restarter or Receiver functionality	BGP neighbor does not support long-lived graceful restart (LLGR) restarter mode completely.

Copyright © 2018, Juniper Networks, Inc.

Table 91: show bgp neighbor Output Fields (continued)

Field Name	Field Description
Peer does not support LLGR Restarter functionality	BGP neighbor does not support long-lived graceful restart (LLGR) restarter mode for any family.
Authentication key change	(Appears only if the **authentication-keychain** statement has been configured) Name of the authentication keychain enabled.
Authentication algorithm	(Appears only if the **authentication-algorithm** statement has been configured) Type of authentication algorithm enabled: **hmac** or **md5**.
Address families configured	Names of configured address families for the VPN.
BGP-Static Advertisement Policy	Name of the BGP static policy that is configured on the peer.
Local Address	Address of the local routing device.
Remove-private options	Options associated with the **remove-private** statement.
Holdtime	Hold time configured with the **hold-time** statement. The hold time is three times the interval at which keepalive messages are sent.
Flags for NLRI inet-label-unicast	Flags related to labeled-unicast: • **TrafficStatistics**—Collection of statistics for labeled-unicast traffic is enabled.
Traffic statistics	Information about labeled-unicast traffic statistics: • **Options**—Options configured for collecting statistics about labeled-unicast traffic. • **File**—Name and location of statistics log files. • **size**—Size of all the log files, in bytes. • **files**—Number of log files.
Traffic Statistics Interval	Time between sample periods for labeled-unicast traffic statistics, in seconds.
Preference	Preference value configured with the **preference** statement.
Outbound Timer	Time for which the route is available in Junos OS routing table before it is exported to BGP. This field is displayed in the output only if the **out-delay** parameter is configured to a non-zero value.
Number of flaps	Number of times the BGP session has gone down and then come back up.
Peer ID	Router identifier of the peer.
Group index	Index number for the BGP peer group. The index number differentiates between groups when a single BGP group is split because of different configuration options at the group and peer levels.

Copyright © 2018, Juniper Networks, Inc.

Table 91: show bgp neighbor Output Fields (continued)

Field Name	Field Description
Peer index	Index that is unique within the BGP group to which the peer belongs.
Local ID	Router identifier of the local routing device.
Local Interface	Name of the interface on the local routing device.
Active holdtime	Hold time that the local routing device negotiated with the peer.
Keepalive Interval	Keepalive interval, in seconds.
BFD	Status of BFD failure detection.
Local Address	Name of directly connected interface over which direct EBGP peering is established.
NLRI and times for LLGR configured on peer	Names of address families and stale time for BGP long-lived graceful restart configured on the BGP peer or neighbor. Times are displayed using the routing protocol daemon (rpd) %#0T format: **<weeks>w<days>d <hours>:<minutes>:<seconds>** Zero leading elements are omitted, for example, a value less than one week do not include the weeks.
NLRI and times that peer supports LLGR Restarter for	Names of address families and stale time that the BGP peer supports for restarter mode for BGP long-lived graceful restart. Times are displayed using the routing protocol daemon (rpd) %#0T format: **<weeks>w<days>d <hours>:<minutes>:<seconds>** Zero leading elements are omitted, for example, a value less than one week do not include the weeks.
NLRI that peer saved LLGR forwarding for	Name of the address family for which the BGP peer saved BGP long-lived graceful restart forwarding.
Graceful Restart Details	Amount of time that is remaining until LLGR expires and the time remaining on the GR stale timer, along with RIB details, are displayed while LLGR receiver mode is active (a peer that negotiated LLGR has disconnected and not yet reconnected).
NLRI we are holding stale routes for	Names of address families (NLRIs) for which that stale routes are held or preserved when BGP graceful restart receiver mode is active for a neighbor.
Time until end-of-rib is assumed for stale routes	Amount of time remaining on the stale timer until which end-of-RIB (EoR) markers are assumed when BGP graceful restart receiver mode is active for a neighbor. Time is displayed in Coordinated Universal Time (UTC) format (YYYY-MM-DD-HH:MM:SS). Note that the stale timer display (‘Time until end-of-rib is assumed’) is also present when a session is active, but the neighbor as not yet sent all of the end-of-rib indications.
Time until stale routes are deleted or become long-lived stale	Amount of time up to which stale routes are deleted or become long-lived stale routes when BGP graceful restart receiver mode is active for a neighbor.

 Copyright © 2018, Juniper Networks, Inc.

Table 91: show bgp neighbor Output Fields (continued)

Field Name	Field Description
NLRI for restart configured on peer	Names of address families configured for restart.
NLRI advertised by peer	Address families supported by the peer: **unicast** or **multicast**.
NLRI for this session	Address families being used for this session.
Peer supports Refresh capability	Remote peer's ability to send and request full routing table readvertisement (route refresh capability). For more information, see RFC 2918, *Route Refresh Capability for BGP-4*.
Restart time configured on peer	Configured time allowed for restart on the neighbor.
Stale routes from peer are kept for	When graceful restart is negotiated, the maximum time allowed to hold routes from neighbors after the BGP session has gone down.
Peer does not support Restarter functionality	Graceful restart restarter-mode is disabled on the peer.
Peer does not support Receiver functionality	Graceful restart helper-mode is disabled on the peer.
Restart time requested by this peer	Restart time requested by this neighbor during capability negotiation.
Restart flag received from the peer	When this field appears, the BGP speaker has restarted (Restarting), and this peer should not wait for the **end-of-rib** marker from the speaker before advertising routing information to the speaker.
NLRI that peer supports restart for	Neighbor supports graceful restart for this address family.
NLRI peer can save forwarding state	Neighbor supporting this address family saves all forwarding states.
NLRI that peer saved forwarding for	Neighbor saves all forwarding states for this address family.
NLRI that restart is negotiated for	Router supports graceful restart for this address family.
NLRI of received end-of-rib markers	Address families for which end-of-routing-table markers are received from the neighbor.
NLRI of all end-of-rib markers sent	Address families for which end-of-routing-table markers are sent to the neighbor.
Peer supports 4 byte AS extension (peer-as 1)	Peer understands 4-byte AS numbers in BGP messages. The peer is running Junos OS Release 9.1 or later.

Copyright © 2018, Juniper Networks, Inc.

Table 91: show bgp neighbor Output Fields (continued)

Field Name	Field Description
NLRIs for which peer can receive multiple paths	Appears in the command output of the local router if the downstream peer is configured to receive multiple BGP routes to a single destination, instead of only receiving the active route. Possible value is inet-unicast.
NLRIs for which peer can send multiple paths: inet-unicast	Appears in the command output of the local router if the upstream peer is configured to send multiple BGP routes to a single destination, instead of only sending the active route. Possible value is inet-unicast.
Table inet.*number*	Information about the routing table: • **RIB State**—BGP is in the graceful restart process for this routing table: **restart is complete** or **restart in progress**. • **Bit**—Number that represents the entry in the routing table for this peer. • **Send state**—State of the BGP group: **in sync**, **not in sync**, or **not advertising**. • **Active prefixes**—Number of prefixes received from the peer that are active in the routing table. • **Received prefixes**—Total number of prefixes from the peer, both active and inactive, that are in the routing table. • **Accepted prefixes**—Total number of prefixes from the peer that have been accepted by a routing policy. • **Suppressed due to damping**—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.
Last traffic (seconds)	Last time any traffic was received from the peer or sent to the peer, and the last time the local routing device checked.
Input messages	Messages that BGP has received from the receive socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.
Output messages	Messages that BGP has written to the transmit socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.
Input dropped path attributes	Information about dropped path attributes: • Code—Path attribute code. • Count—Path attribute count.
Input ignored path attributes	Information about ignored path attributes: • Code—Path attribute code. • Count—Path attribute count.

 Copyright © 2018, Juniper Networks, Inc.

Table 91: show bgp neighbor Output Fields (continued)

Field Name	Field Description
Output queue	Number of BGP packets that are queued to be transmitted to a particular neighbor for a particular routing table. Output queue **0** is for unicast NLRIs, and queue **1** is for multicast NLRIs. It also specifies the routing table name and the NLRI that the table was advertised through, in the format **(*routing table name*, *NLRI*)**. NOTE: The output queue of routing tables that are not advertised, will only show up at **extensive** output level.
Trace options	Configured tracing of BGP protocol packets and operations.
Trace file	Name of the file to receive the output of the tracing operation.
Filter Updates recv	(**orf** option only) Number of outbound-route filters received for each configured address family. NOTE: The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.
Immediate	(**orf** option only) Number of route updates received with the immediate flag set. The immediate flag indicates that the BGP peer should readvertise the updated routes. NOTE: The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.
Filter	(**orf** option only) Type of prefix filter received: **prefix-based** or **extended-community**.
Received filter entries	(**orf** option only) List of received filters displayed.
seq	(**orf** option only) Numerical order assigned to this prefix entry among all the received outbound route filter prefix entries.
prefix	(**orf** option only) Address for the prefix entry that matches the filter.
minlength	(**orf** option only) Minimum prefix length, in bits, required to match this prefix.
maxlength	(**orf** option only) Maximum prefix length, in bits, required to match this prefix.
match	(**orf** option only) For this prefix match, whether to **permit** or **deny** route updates.

Sample Output

show bgp neighbor

`user@host >` **show bgp neighbor**

For M Series, MX Series, and T Series routers running Junos OS Release 16.1 or later, the **show bgp neighbor** output includes the BGP group the peer belongs to, the routing instance (if any) that the peer is configured in, and the routing instance that the peer is using for the forwarding context (if applicable). An example follows.

```
Peer: 10.255.7.250+179 AS 10   Local: 10.255.7.248+63740 AS 10
  Group: toAsbr2               Routing-Instance: master
  Forwarding routing-instance: toAs2
    Type: Internal    State: Established    Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ redist_static ]
  Options: <Preference LocalAddress PeerAS Refresh>
  Options: <AdvertiseBGPStatic>
  Local Address: 10.255.7.248 Holdtime: 90 Preference: 170 Outbound Timer: 50
  Number of flaps: 0
  Peer ID: 10.255.7.250    Local ID: 10.255.7.248      Active Holdtime: 90
  Keepalive Interval: 30         Group index: 0    Peer index: 0
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer supports 4 byte AS extension (peer-as 10)
  Peer does not support Addpath
  NLRI that we support extended nexthop encoding for: inet-unicast
  NLRI that peer supports extended nexthop encoding for: inet-unicast

  Table inet.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:              1
    Received prefixes:            1
    Accepted prefixes:            1
    Suppressed due to damping:    0
    Advertised prefixes:          1
  Last traffic (seconds): Received 9    Sent 5    Checked 5
  Input messages:  Total 36     Updates 2       Refreshes 0     Octets 718
  Output messages: Total 37     Updates 1       Refreshes 0     Octets 796
  Output Queue[0]: 0  (inet.0, inet-unicast)

Peer: 10.255.162.214+52193 AS 100 Local: 10.255.167.205+179 AS 100
   Type: Internal    State: Established  (route reflector client)Flags: <Sync>
   Last State: OpenConfirm   Last Event: RecvKeepAlive
   Last Error: None
   Options: <Preference LocalAddress Cluster AddressFamily Rib-group Refresh>
   Address families configured: inet-unicast inet-vpn-unicast route-target
   Local Address: 10.255.167.205 Holdtime: 90 Preference: 170
   Number of flaps: 0
   Peer ID: 10.255.162.214  Local ID: 10.255.167.205    Active Holdtime: 90
   Keepalive Interval: 30         Group index: 0    Peer index: 1
```

show bgp neighbor (CLNS)

```
user@host> show bgp neighbor
Peer: 10.245.245.1+179 AS 200  Local: 10.245.245.3+3770 AS 100
  Type: External    State: Established    Flags: <ImportEval Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Multihop Preference LocalAddress HoldTime AddressFamily PeerAS
Rib-group Refresh>
```

```
  Address families configured: iso-vpn-unicast
  Local Address: 10.245.245.3 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.245.245.1     Local ID: 10.245.245.3     Active Holdtime: 90
  Keepalive Interval: 30         Peer index: 0
  NLRI advertised by peer: iso-vpn-unicast
  NLRI for this session: iso-vpn-unicast
  Peer supports Refresh capability (2)
  Table bgp.isovpn.0 Bit: 10000
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: in sync
    Active prefixes:              3
    Received prefixes:            3
    Suppressed due to damping:    0
    Advertised prefixes:          3
  Table aaaa.iso.0
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: not advertising
    Active prefixes:              3
    Received prefixes:            3
    Suppressed due to damping:    0
  Last traffic (seconds): Received 6    Sent 5    Checked 5
  Input messages:  Total 1736   Updates 4       Refreshes 0     Octets 33385
  Output messages: Total 1738   Updates 3       Refreshes 0     Octets 33305
  Output Queue[0]: 0  (bgp.isovpn.0, iso-vpn-unicast)
  Output Queue[1]: 0  (aaaa.iso.0, iso-vpn-unicast)
```

show bgp neighbor (Layer 2 VPN)

```
user@host> show bgp neighbor
Peer: 10.69.103.2     AS 65536 Local: 10.69.103.1     AS 65539
  Type: External    State: Active         Flags: <ImportEval>
  Last State: Idle          Last Event: Start
  Last Error: None
  Export: [ BGP-INET-import ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily   PeerAS
 Refresh>
  Address families configured: inet-unicast
  Local Address: 10.69.103.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
Peer: 10.69.104.2     AS 65539 Local: 10.69.104.1     AS 65539
  Type: External    State: Active         Flags: <ImportEval>
  Last State: Idle          Last Event: Start
  Last Error: None
  Export: [ BGP-L-import ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
 Refresh>
  Address families configured: inet-labeled-unicast
  Local Address: 10.69.104.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
Peer: 10.255.14.182+179 AS 69    Local: 10.255.14.176+2131 AS 69
  Type: Internal    State: Established    Flags: <ImportEval>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
  Rib-group Refresh>
  Address families configured: inet-vpn-unicast l2vpn
  Local Address: 10.255.14.176 Holdtime: 90 Preference: 170
  Number of flaps: 0
```

```
Peer ID: 10.255.14.182    Local ID: 10.255.14.176    Active Holdtime: 90
Keepalive Interval: 30
NLRI for restart configured on peer: inet-vpn-unicast l2vpn
NLRI advertised by peer: inet-vpn-unicast l2vpn
NLRI for this session: inet-vpn-unicast l2vpn
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-vpn-unicast l2vpn
NLRI peer can save forwarding state: inet-vpn-unicast l2vpn
NLRI that peer saved forwarding for: inet-vpn-unicast l2vpn
NLRI that restart is negotiated for: inet-vpn-unicast l2vpn
NLRI of received end-of-rib markers: inet-vpn-unicast l2vpn
Table bgp.l3vpn.0 Bit: 10000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:           10
  Received prefixes:         10
  Suppressed due to damping:  0
Table bgp.l2vpn.0 Bit: 20000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:            1
  Received prefixes:          1
  Suppressed due to damping:  0
Table BGP-INET.inet.0 Bit: 30000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:            2
  Received prefixes:          2
  Suppressed due to damping:  0
Table BGP-L.inet.0 Bit: 40000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:            2
  Received prefixes:          2
  Suppressed due to damping:  0
Table LDP.inet.0 Bit: 50000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:            1
  Received prefixes:          1
  Suppressed due to damping:  0
Table OSPF.inet.0 Bit: 60000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:            2
  Received prefixes:          2
  Suppressed due to damping:  0
Table RIP.inet.0 Bit: 70000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:            2
```

 Copyright © 2018, Juniper Networks, Inc.

```
    Received prefixes:          2
    Suppressed due to damping:  0
  Table STATIC.inet.0 Bit: 80000
    RIB State: BGP restart is complete
    RIB State: VPN restart in progress
    Send state: in sync
    Active prefixes:            1
    Received prefixes:          1
    Suppressed due to damping:  0
  Table L2VPN.l2vpn.0 Bit: 90000
    RIB State: BGP restart is complete
    RIB State: VPN restart in progress
    Send state: in sync
    Active prefixes:            1
    Received prefixes:          1
    Suppressed due to damping:  0
  Last traffic (seconds): Received 0    Sent 0    Checked 0
  Input messages:  Total 14     Updates 13      Refreshes 0     Octets 1053
  Output messages: Total 3      Updates 0       Refreshes 0     Octets 105
  Output Queue[0]: 0  (bgp.l3vpn.0, inet-vpn-unicast)
  Output Queue[1]: 0  (bgp.l2vpn.0, inet-vpn-unicast)
  Output Queue[2]: 0  (BGP-INET.inet.0, inet-vpn-unicast)
  Output Queue[3]: 0  (BGP-L.inet.0, inet-vpn-unicast)
  Output Queue[4]: 0  (LDP.inet.0, inet-vpn-unicast)
  Output Queue[5]: 0  (OSPF.inet.0, inet-vpn-unicast)
  Output Queue[6]: 0  (RIP.inet.0, inet-vpn-unicast)
  Output Queue[7]: 0  (STATIC.inet.0, inet-vpn-unicast)
  Output Queue[8]: 0  (L2VPN.l2vpn.0, inet-vpn-unicast)
```

show bgp neighbor (Layer 3 VPN) (Not supported on the OCX Series.)

```
user@host> show bgp neighbor
Peer: 192.0.2.0.179     AS 10045 Local: 192.0.2.1+1214    AS 10045
  Type: Internal     State: Established     Flags: <ImportEval>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ match-all ] Import: [ match-all ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
   Rib-group Refresh>
  Address families configured: inet-vpn-unicast
  Local Address: 192.0.2.1 Holdtime: 90 Preference: 170
  Flags for NLRI inet-labeled-unicast: TrafficStatistics
  Traffic Statistics: Options: all File: /var/log/bstat.log
                                       size 131072 files 10
  Traffic Statistics Interval: 60
  Number of flaps: 0
  Peer ID: 192.168.1.110    Local ID: 192.168.1.111    Active Holdtime: 90
  Keepalive Interval: 30
  NLRI for restart configured on peer: inet-vpn-unicast
  NLRI advertised by peer: inet-vpn-unicast
  NLRI for this session: inet-vpn-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-vpn-unicast
  NLRI peer can save forwarding state: inet-vpn-unicast
  NLRI that peer saved forwarding for: inet-vpn-unicast
  NLRI that restart is negotiated for: inet-vpn-unicast
  NLRI of received end-of-rib markers: inet-vpn-unicast
  NLRI of all end-of-rib markers sent: inet-vpn-unicast
```

```
  Table bgp.l3vpn.0 Bit: 10000
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: in sync
    Active prefixes:            2
    Received prefixes:          2
    Suppressed due to damping:  0
  Table vpn-green.inet.0 Bit: 20001
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: in sync
    Active prefixes:            2
    Received prefixes:          2
    Suppressed due to damping:  0
  Last traffic (seconds): Received 15   Sent 20   Checked 20
  Input messages:  Total 40     Updates 2       Refreshes 0     Octets 856
  Output messages: Total 44     Updates 2       Refreshes 0     Octets 1066
  Output Queue[0]: 0  (bgp.l3vpn.0, inet-vpn-unicast)
  Output Queue[1]: 0  (vpn-green.inet.0, inet-vpn-unicast)
  Trace options: detail packets
  Trace file: /var/log/bgpgr.log size 131072 files 10
```

show bgp neighbor neighbor-address

```
user@host> show bgp neighbor 192.168.1.111
Peer: 10.255.245.12+179 AS 35  Local: 10.255.245.13+2884 AS 35
  Type: Internal    State: Established  (route reflector client)Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress HoldTime Cluster AddressFamily Rib-group
  Refresh>
  Options: RFC6514CompliantSafi129
  Address families configured: inet-vpn-unicast inet-labeled-unicast
  Local Address: 10.255.245.13 Holdtime: 90 Preference: 170
  Flags for NLRI inet-vpn-unicast: AggregateLabel
  Flags for NLRI inet-labeled-unicast: AggregateLabel
  Number of flaps: 0
  Peer ID: 10.255.245.12    Local ID: 10.255.245.13    Active Holdtime: 90
  Keepalive Interval: 30
BFD: disabled
  NLRI advertised by peer: inet-vpn-unicast inet-labeled-unicast
  NLRI for this session: inet-vpn-unicast inet-labeled-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 300
  Stale routes from peer are kept for: 60
  Restart time requested by this peer: 300
  NLRI that peer supports restart for: inet-unicast inet6-unicast
  NLRI that restart is negotiated for: inet-unicast inet6-unicast
  NLRI of received end-of-rib markers: inet-unicast inet6-unicast
  NLRI of all end-of-rib markers sent: inet-unicast inet6-unicast
  Table inet.0 Bit: 10000
    RIB State: restart is complete
    Send state: in sync
    Active prefixes: 4
    Received prefixes: 6
    Suppressed due to damping: 0
  Table inet6.0 Bit: 20000
    RIB State: restart is complete
    Send state: in sync
    Active prefixes: 0
    Received prefixes: 2
```


Copyright © 2018, Juniper Networks, Inc.

```
    Suppressed due to damping: 0
  Last traffic (seconds): Received 3    Sent 3    Checked 3
  Input messages:  Total 9      Updates 6       Refreshes 0     Octets 403
  Output messages: Total 7      Updates 3       Refreshes 0     Octets 365
  Output Queue[0]: 0  (inet.0, inet-unicast)
  Output Queue[1]: 0  (inet6.0, inet6-unicast)
  Trace options: detail packets
  Trace file: /var/log/bgpgr size 131072 files 10
```

show bgp neighbor neighbor-address

```
user@host> show bgp neighbor 192.168.4.222
Peer: 192.168.4.222+4902 AS 65501 Local: 192.168.4.221+179 AS 65500
  Type: External    State: Established    Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: Cease
  Export: [ export-policy ] Import: [ import-policy ]
  Options: <Preference HoldTime AddressFamily PeerAS PrefixLimit Refresh>
  Address families configured: inet-unicast inet-multicast
  Holdtime: 60000 Preference: 170
  Number of flaps: 4
  Last flap event: RecvUpdate
  Error: 'Cease' Sent: 5 Recv: 0
  Peer ID: 10.255.245.6     Local ID: 10.255.245.5     Active Holdtime: 60000
  Keepalive Interval: 20000      Peer index: 0
  BFD: disabled, down
  Local Interface: fxp0.0
  NLRI advertised by peer: inet-unicast inet-multicast
  NLRI for this session: inet-unicast inet-multicast
  Peer supports Refresh capability (2)
  Table inet.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:              8
    Received prefixes:            10
    Accepted prefixes:            10
    Suppressed due to damping:    0
    Advertised prefixes:          3
  Table inet.2 Bit: 20000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:              0
    Received prefixes:            0
    Accepted prefixes:            0
    Suppressed due to damping:    0
    Advertised prefixes:          0
  Last traffic (seconds): Received 357  Sent 357  Checked 357
  Input messages:  Total 4 Updates 2  Refreshes 0  Octets 211
  Output messages: Total 4 Updates 1  Refreshes 0  Octets 147
  Output Queue[0]: 0  (inet.0, inet-unicast)
  Output Queue[1]: 0  (inet.2, inet-multiicast)
  Trace options:  all
  Trace file: /var/log/bgp size 10485760 files 10
```

show bgp neighbor neighbor-address (BGP Graceful Restart Enabled)

```
user@router> show bgp neighbor 10.255.255.16

  Peer: 10.255.255.16 AS 100     Local: 10.255.255.12 AS 100
    Type: Internal    State: Active         Flags: <>
```

```
    Last State: Idle          Last Event: Start
    Last Error: None
    Options: <Preference LocalAddress AddressFamily Rib-group Refresh>
    Options: <LLGR>
    Address families configured: l2vpn
    Local Address: 10.255.255.12 Holdtime: 90 Preference: 170
    NLRI l2vpn:
    Number of flaps: 6
    Last flap event: Restart
    NLRI we are holding stale routes for: inet-vpn-unicast
    Time until stale routes are deleted or become long-lived stale: 00:01:57
    Time until end-of-rib is assumed for stale routes: 00:04:43
    Table bgp.l3vpn.0
      RIB State: BGP restart is complete
      RIB State: VPN restart is complete
      Send state: not advertising
      Active prefixes:              0
      Received prefixes:            7
      Accepted prefixes:            7
      Suppressed due to damping:    0
    Table foo.inet.0 Bit: 30000
      RIB State: BGP restart is complete
      RIB State: VPN restart is complete
      Send state: not in sync
      Active prefixes:              0
      Received prefixes:            7
      Accepted prefixes:            7
      Suppressed due to damping:    0
```

show bgp neighbor neighbor-address (BGP Long-Lived Graceful Restart)

user@router> **show bgp neighbor 10.4.12.11**

```
  Peer: 10.4.12.11 AS 100        Local: 10.6.128.225 AS 100
    Type: Internal    State: Active         Flags: <>
    Last State: Idle          Last Event: Start
    Last Error: None
    Export: [ foo ]
    Options: <Preference LocalAddress Refresh GracefulRestart>
    Options: <LLGR>
    Local Address: 10.6.128.225 Holdtime: 90 Preference: 170
    Number of flaps: 3
    Last flap event: Restart
    Error: 'Cease' Sent: 0 Recv: 1
    Time until long-lived stale routes deleted: inet-vpn-unicast 10:00:22
route-target 10:00:22
    Table bgp.l3vpn.0
      RIB State: BGP restart is complete
      RIB State: VPN restart is complete
      Send state: not advertising
      Active prefixes:              0
      Received prefixes:            7
      Accepted prefixes:            7
      Suppressed due to damping:    0
    Table foo.inet.0 Bit: 30000
      RIB State: BGP restart is complete
      RIB State: VPN restart is complete
      Send state: not in sync
      Active prefixes:              0
      Received prefixes:            7
```


Copyright © 2018, Juniper Networks, Inc.

```
      Accepted prefixes:             7
      Suppressed due to damping:     0
```

show bgp neighbor orf neighbor-address detail

```
user@host > show bgp neighbor orf 192.168.165.56 detail
Peer: 192.168.165.56+179 Type: External
  Group: ext1

  inet-unicast
    Filter updates recv:           1 Immediate:          1
    Filter: prefix-based  receive
      Received filter entries:
        seq 1: prefix 2.2.2.2/32: minlen 32: maxlen 32: match deny:

  inet6-unicast
    Filter updates recv:           0 Immediate:          1
    Filter: prefix-based  receive
      Received filter entries:
        *:*
```

show bgp neighbor logical-system

```
user@host > show bgp neighbor logical-system ITR1
Peer: 10.79.8.2+179 AS 65536   Local: 10.79.8.1+50891 AS 65500
  Description: MX1
  Type: External    State: Established    Flags: <ImportEval Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
 ....
  Table inet.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:               1
    Received prefixes:             1
    Accepted prefixes:             1
    Suppressed due to damping:     0
    Advertised prefixes:           10
    Stale prefixes:                4: <=new, line only appears if count is non-0
It is the Number of prefixes marked as stale;
    LLGR-stale prefixes:           5: <=new, line only appears if count is non-0
It is the Number of prefixes marked as LLGR-stale
```

show bgp neighbor output-queue

```
user@host > show bgp neighbor output-queue
Peer: 192.0.2.2+179 AS 103       Local: 192.0.2.1+50799 AS 102
  Output Queue[0]: 0            (inet.0, inet-unicast)
    Priority 1 :   0
    Priority 2 :   0
    Priority 3 :   0
    Priority 4 :   0
    Priority 5 :   0
    Priority 6 :   0
    Priority 7 :   0
    Priority 8 :   0
    Priority 9 :   0
    Priority 10:   0
    Priority 11:   0
    Priority 12:   0
```

```
      Priority 13:   0
      Priority 14:   0
      Priority 15:   0
      Priority 16:   0
      Expedited  :   0
```

show bgp neighbor (Segment Routing Traffic Engineering)

```
user@host > show bgp neighbor
run show bgp neighbor 1.1.1.254
   Peer: 1.1.1.254+60180 AS 100   Local: 1.1.1.1+179 AS 100
  Group: toB                   Routing-Instance: master
  Forwarding routing-instance: master
  Type: Internal    State: Established    Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress>
  Address families configured: inet-segment-routing-te
  Local Address: 1.1.1.1 Holdtime: 90 Preference: 170 Local AS: 100 Local System
 AS: 0
  Number of flaps: 0
  Peer ID: 128.9.150.15    Local ID: 128.9.150.110     Active Holdtime: 90
  Keepalive Interval: 30         Group index: 0    Peer index: 0
  I/O Session Thread: bgpio-0 State: Enabled
  BFD: disabled, down
  NLRI for restart configured on peer: inet-segment-routing-te
  NLRI advertised by peer: inet-segment-routing-te
  NLRI for this session: inet-segment-routing-te
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  Restart flag received from the peer: Notification
  NLRI that restart is negotiated for: inet-segment-routing-te
  Peer does not support LLGR Restarter functionality
  Peer supports 4 byte AS extension (peer-as 100)
  Peer does not support Addpath
  Last traffic (seconds): Received 17628 Sent 25  Checked 17628
  Input messages:  Total 2      Updates 0       Refreshes 0     Octets 82
  Output messages: Total 1      Updates 0       Refreshes 0     Octets 19
  Trace options:  all
  Trace file: /var/log/bgp.log size 10485760 files 10
```

 Copyright © 2018, Juniper Networks, Inc.

show bgp replication

Syntax **show bgp replication**

Release Information Command introduced in Junos OS Release 8.5.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Support for **logical-system** option introduced in Junos OS Release 13.3

Description Displays the status of BGP state replication between the master and backup Routing Engines on devices that have nonstop active routing configured on them.

CAUTION: Before attempting nonstop active routing switchover, check the output of **show bgp replication** to confirm that BGP routing table synchronization has completed on the backup Routing Engine. The **complete** status in the output of **show task replication** only indicates that the socket replication has completed and the BGP synchronization is in progress.

To determine whether BGP synchronization is complete, you must check the **Protocol state** and **Synchronization state** fields in the output of **show bgp replication** on the master Routing Engine. The **Protocol state** must be **idle** and the **Synchronization state** must be **complete**. If you perform NSR switchover before the BGP synchronization has completed, the BGP session might flap.

Options This command has no options.

Required Privilege Level view

Related Documentation
- *show bgp replication logical-system*

List of Sample Output
show bgp replication (for Master) on page 1703
show bgp replication (for Master) on page 1703
show bgp replication (for Backup) on page 1703

Output Fields Table 92 on page 1702 lists the output fields for the **show bgp replication** command. Output fields are listed in the approximate order in which they appear.

Table 92: show bgp replication Output Fields

Field Name	Field Description
Precision timer registration	State of BGP precision timer feature in the kernel. • **Registered**—BGP registers with the precision-timer feature in the kernel for auto keepalive generation after switchover. • **NotRegistered**—Keepalive format of BGP is not registered.
session state	State of the current internal BGP state replication session, Up or Down, and the duration for which the session has been in the indicated state.
flaps	Total number of flaps that occurred.
protocol state	Current state of the protocol operation, Active, Connect, Idle, and the duration for which the protocol has been in the indicated state.
synchronization state	Synchronization state at the time of executing the command. The states can be: • **Idle** • **Neighbor**—Indicates that the neighbor state synchronization is in progress. • **AckWait**—Indicates that the request processing is over. • **ORF**—Indicates that the outbound routing filter synchronization is in progress. • **RIB**—Indicates that the routing table synchronization is in progress. • **Complete**
number of peers waiting	Total number of peers waiting for various messages: • **AckWait**—Number of peers waiting for a connection establishment or completed acknowledgment messages. • **SoWait**—Number of peers waiting for TCP socket-related operations. • **Scheduled**—Number of peers being synchronized.
messages sent	Number of various types of messages that have been sent since internal replication session became active: • **Open**—Number of Open messages sent. • **Establish**—Number of connection establishment acknowledgment messages sent. • **Update**—Number of update messages sent. • **Error**—Number of error messages sent. • **Complete**—Number of connection complete acknowledgment messages sent.

messages received	Total number of messages received:
	• **Open**—Number of Open messages received.
	• **Request**—Number of request messages received:
	  ◦ **Wildcard**—Number of requests received that used wildcards in the target address.
	  ◦ **Targeted**—Number of requests received that used a specific address.
	• **EstablishAck**—Number of connection establishment acknowledgement messages received.
	• **CompleteAck**—Number of connection completed acknowledgement messages received.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show bgp replication (for Master)

```
user@host> show bgp replication
Synchronization master:
  Precision timer registration: Registered
  Session state: Up, Since: 10:14
  Flaps: 1, Last flap reason: Backup closed connection
  Protocol state: Idle, Since: 10:14
  Synchronization state: Complete
  Number of peers waiting: AckWait: 0, SoWait: 0, Scheduled: 0
  Messages sent: Open 1, Establish 11, GrHelper 0, Update 0, GrStaleLabel 0 Error
 0, Complete 1
  Messages received: Open 1, Request 1 wildcard 0 targeted, EstablishAck 11,
GrHelperAck 0, CompleteAck 1
```

show bgp replication (for Master)

```
user@host> show bgp replication
Synchronization master:
  Session state: Up, Since: 44:07
  Flaps: 0
  Protocol state: Idle, Since: 14
  Synchronization state: Complete
  Number of peers waiting: AckWait: 0, SoWait: 0, Scheduled: 0
  Messages sent: Open 1, Establish 924, Update 381, Error 60, Complete 114
  Messages received: Open 1, Request 1 wildcard 113 targeted, EstablishAck 924,
CompleteAck 114
```

show bgp replication (for Backup)

```
user@host> show bgp replication
Synchronization backup:
  State: Established 13 ago
  , Unsync timer: 2

    Unsync entry queue:
        Instance: 0 Neighbor: 30.30.30.1 elapsed: 7
        Instance: 0 Neighbor: 40.40.40.3 elapsed: 7
        Instance: 0 Neighbor: 40.40.40.4 elapsed: 7
        Instance: 0 Neighbor: 40.40.40.5 elapsed: 7
        Instance: 0 Neighbor: 40.40.40.6 elapsed: 7
        Instance: 0 Neighbor: 40.40.40.1 elapsed: 7
        Instance: 0 Neighbor: 40.40.40.2 elapsed: 7
```

Copyright © 2018, Juniper Networks, Inc.

show bgp summary

List of Syntax

Syntax on page 1704
Syntax (EX Series Switch and QFX Series) on page 1704

Syntax

show bgp summary
<exact-instance *instance-name*>
<group *group-name*>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and QFX Series)

show bgp summary
<exact-instance *instance-name*>
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
exact-instance option introduced in Junos OS Release 11.4.
group option introduced in Junos OS Release 13.3.

Description

Display BGP summary information.

Options

none—Display BGP summary information for all routing instances.

exact-instance ***instance-name***—(Optional) Display information for the specified instance only.

group—Display overview of bgp information for a particular group

instance ***instance-name***—(Optional) Display information for all routing instances whose name begins with this string (for example, **cust1**, **cust11**, and **cust111** are all displayed when you run the **show bgp summary instance cust1** command). The instance name can be master for the main instance, or any valid configured instance name or its prefix.

logical-system (all | ***logical-system-name*****)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show bgp summary (When a Peer Is Not Established) on page 1707
show bgp summary (When a Peer Is Established) on page 1707
show bgp summary (CLNS) on page 1707
show bgp summary (Layer 2 VPN) on page 1708
show bgp summary (Layer 3 VPN) on page 1708

 Copyright © 2018, Juniper Networks, Inc.

show bgp summary group on page 1708
show bgp summary (BGP Graceful Restart or Long-Lived Graceful Restart) on page 1709

Output Fields Table 93 on page 1705 describes the output fields for the **show bgp summary** command. Output fields are listed in the approximate order in which they appear.

Table 93: show bgp summary Output Fields

Field Name	Field Description
Groups	Number of BGP groups.
Peers	Number of BGP peers.
Down peers	Number of down BGP peers.
Table	Name of routing table.
Tot Paths	Total number of paths.
Act Paths	Number of active routes.
Suppressed	Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.
History	Number of withdrawn routes stored locally to keep track of damping history.
Damp State	Number of routes with a figure of merit greater than zero, but still active because the value has not reached the threshold at which suppression occurs.
Pending	Routes in process by BGP import policy.
Peer	Address of each BGP peer. Each peer has one line of output.
AS	Peer's AS number.
InPkt	Number of packets received from the peer.
OutPkt	Number of packets sent to the peer.
OutQ	Number of BGP packets that are queued to be transmitted to a particular neighbor. It normally is **0** because the queue usually is emptied quickly.
Flaps	Number of times the BGP session has gone down and then come back up.
Last Up/Down	Last time since the neighbor transitioned to or from the established state.

Copyright © 2018, Juniper Networks, Inc.

Table 93: show bgp summary Output Fields (continued)

Field Name	Field Description
State|#Active /Received/Accepted /Damped	Multipurpose field that displays information about BGP peer sessions. The field’s contents depend upon whether a session is established and whether it was established on the main routing device or in a routing instance. • If a peer is not established, the field shows the state of the peer session: **Active**, **Connect**, or **Idle**. In general, the Idle state is the first stage of a connection. BGP is waiting for a Start event. A session can be idle for other reasons as well. The reason that a session is idle is sometimes displayed. For example: **Idle (Removal in progress)** or **Idle (LicenseFailure)**. • If a BGP session is established on the main routing device, the field shows the number of active, received, accepted, and damped routes that are received from a neighbor and appear in the **inet.0** (main) and **inet.2** (multicast) routing tables. For example, **8/10/10/2 and 2/4/4/0** indicate the following: ◦ 8 active routes, 10 received routes, 10 accepted routes, and 2 damped routes from a BGP peer appear in the **inet.0** routing table. ◦ 2 active routes, 4 received routes, 4 accepted routes, and no damped routes from a BGP peer appear in the **inet.2** routing table. • If a BGP session is established in a routing instance, the field indicates the established **(Establ)** state, identifies the specific routing table that receives BGP updates, and shows the number of active, received, and damped routes that are received from a neighbor. For example, **Establ VPN-AB.inet.0: 2/4/0** indicates the following: ◦ The BGP session is established. ◦ Routes are received in the **VPN-AB.inet.0** routing table. ◦ The local routing device has two active routes, four received routes, and no damped routes from a BGP peer. When a BGP session is established, the peers are exchanging update messages. NOTE: When graceful restart or LLGR helper mode is active, the RIB information is now displayed by the **show bgp summary** command. If a BGP session is established on the main routing device, the field shows the number of active, received, accepted, and damped routes that are received from a neighbor and appear in the inet.0 (main) and inet.2 (multicast) routing tables. For example, 8/10/10/2 and 2/4/4/0 indicate the following: • 8 active routes, 10 received routes, 10 accepted routes, and 2 damped routes from a BGP peer appear in the inet.0 routing table. • 2 active routes, 4 received routes, 4 accepted routes, and no damped routes from a BGP peer appear in the inet.2 routing table.

 Copyright © 2018, Juniper Networks, Inc.

Sample Output

show bgp summary (When a Peer Is Not Established)

```
user@host> show bgp summary
Groups: 2 Peers: 4 Down peers: 1
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0                 6          4          0          0          0          0
Peer               AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn 
State|#Active/Received/Damped...
10.0.0.3        65002         86         90       0       2       42:54 0/0/0     

0/0/0
10.0.0.4        65002         90         91       0       1       42:54 0/2/0     

0/0/0
10.0.0.6        65002         87         90       0       3           3 Active
10.1.12.1       65001         89         89       0       1       42:54 4/4/0     

0/0/0
```

show bgp summary (When a Peer Is Established)

```
user@host> show bgp summary
Groups: 1 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0                 6          4          0          0          0          0
Peer               AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn 
State|#Active/Received/Damped...
10.0.0.2        65002      88675      88652       0       2       42:38 2/4/0     

0/0/0
10.0.0.3        65002      54528      54532       0       1     2w4d22h 0/0/0     

 0/0/0
10.0.0.4        65002      51597      51584       0       0     2w3d22h 2/2/0     

0/0/0

user@host> show bgp summary logical-system R3
Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
bgp.l3vpn.0
                       2          2          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn 
State|#Active/Received/Accepted/Damped...
1.1.1.2                   2        204        206       0       0     1:30:59 
Establ
  bgp.l3vpn.0: 2/2/2/0
  red.inet.0: 2/2/2/0
10.1.1.10                 3        206        207       0       0     1:31:36 
Establ
  red.inet.0: 2/2/2/0
```

show bgp summary (CLNS)

```
user@host> show bgp summary
Groups: 1 Peers: 1 Down peers: 0
Peer               AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn 
State|#Active/Received/Damped...
```

Copyright © 2018, Juniper Networks, Inc.

```
10.245.245.1      200       1735       1737        0        0    14:26:12 Establ
  bgp.isovpn.0: 3/3/0
  aaaa.iso.0: 3/3/0
```

show bgp summary (Layer 2 VPN)

```
user@host> show bgp summary
Groups: 1 Peers: 5 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
bgp.l2vpn.0            1          1          0          0          0          0
inet.0                 0          0          0          0          0          0
Peer               AS      InPkt     OutPkt    OutQ   Flaps Last
Up/Dwn State|#Active/Received/Damped...
10.255.245.35   65299         72         74       0       1       19:00 Establ
  bgp.l2vpn.0: 1/1/0
  frame-vpn.l2vpn.0: 1/1/0
10.255.245.36   65299       2164       2423       0       4       19:50 Establ
  bgp.l2vpn.0: 0/0/0
  frame-vpn.l2vpn.0: 0/0/0
10.255.245.37   65299         36         37       0       4       17:07 Establ
  inet.0: 0/0/0
10.255.245.39   65299        138        168       0       6       53:48 Establ
  bgp.l2vpn.0: 0/0/0
  frame-vpn.l2vpn.0: 0/0/0
10.255.245.69   65299        134        140       0       6       53:42 Establ
  inet.0: 0/0/0
```

show bgp summary (Layer 3 VPN)

```
user@host>  show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State Pending
bgp.l3vpn.0            2          2          0          0          0      0
Peer               AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.39.1.5          2         21         22       0       0        6:26 Establ
  VPN-AB.inet.0: 1/1/0
10.255.71.15       1         19         21       0       0        6:17 Establ
  bgp.l3vpn.0: 2/2/0
  VPN-A.inet.0: 1/1/0
  VPN-AB.inet.0: 2/2/0
  VPN-B.inet.0: 1/1/0
```

show bgp summary group

```
user@host>  show bgp summary group Group2
Groups: 3 Peers: 3 Down peers: 3
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       0          0          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.0.0.1                 56          0          0       0       0          51
Idle

user@host>  show bgp summary logical-system R3 group toR4
Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
bgp.l3vpn.0
                       2          2          0          0          0          0
```

 Copyright © 2018, Juniper Networks, Inc.

```
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.1.1.10                 3        207        207       0       0     1:31:40
Establ
  red.inet.0: 2/2/2/0
```

show bgp summary (BGP Graceful Restart or Long-Lived Graceful Restart)

```
user@router> show route receive-protocol bgp 10.4.12.11 detail
  Groups: 2 Peers: 9 Down peers: 1
  ...
  Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
  10.255.255.16           100          7          6       0       4           4
Idle
    bgp.l3vpn.0: 0/7/7/0
    foo.inet.0: 0/7/7/0
```

Copyright © 2018, Juniper Networks, Inc.

show policy damping

List of Syntax

Syntax on page 1710
Syntax (EX Series Switch and QFX Series) on page 1710

Syntax

show policy damping
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and QFX Series)

show policy damping

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display information about BGP route flap damping parameters.

Options

none—Display information about BGP route flap damping parameters.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information

In the output from this command, figure-of-merit values correlate with the probability of future instability of a routing device. Routes with higher figure-of-merit values are suppressed for longer periods of time. The figure-of-merit value decays exponentially over time. A figure-of-merit value of zero is assigned to each new route. The value is increased each time the route is withdrawn or readvertised, or when one of its path attributes changes.

Required Privilege Level

view

Related Documentation

- “Configuring BGP Flap Damping Parameters” in the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*
- clear bgp damping on page 1664
- show route damping on page 2215

List of Sample Output

show policy damping on page 1711

Output Fields

Table 94 on page 1711 describes the output fields for the **show policy damping** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2018, Juniper Networks, Inc.

Table 94: show policy damping Output Fields

Field Name	Field Description
Halflife	Decay half-life, in minutes. The value represents the period during which the accumulated figure-of-merit value is reduced by half if the route remains stable. If a route has flapped, but then becomes stable, the figure-of-merit value for the route decays exponentially. For example, for a route with a figure-of-merit value of 1500, if no incidents occur, its figure-of-merit value is reduced to 750 after 15 minutes and to 375 after another 15 minutes.
Reuse merit	Figure-of-merit value below which a suppressed route can be used again. A suppressed route becomes reusable when its figure-of-merit value decays to a value below a reuse threshold, and the route once again is considered usable and can be installed in the forwarding table and exported from the routing table.
Suppress/cutoff merit	Figure-of-merit value above which a route is suppressed for use or inclusion in advertisements. When a route's figure-of-merit value reaches a particular level, called the cutoff or suppression threshold, the route is suppressed. When a route is suppressed, the routing table no longer installs the route into the forwarding table and no longer exports this route to any of the routing protocols.
Maximum suppress time	Maximum hold-down time, in minutes. The value represents the maximum time that a route can be suppressed no matter how unstable it has been before this period of stability.
Computed values	• **Merit ceiling**—Maximum merit that a flapping route can collect. • **Maximum decay**—Maximum decay half-life, in minutes.

Sample Output

show policy damping

```
user@host> show policy damping
Default damping information:
  Halflife: 15 minutes
  Reuse merit: 750 Suppress/cutoff merit: 3000
  Maximum suppress time: 60 minutes
  Computed values:
    Merit ceiling: 12110
    Maximum decay: 6193
Damping information for "standard-damping":
  Halflife: 10 minutes
  Reuse merit: 4000 Suppress/cutoff merit: 8000
  Maximum suppress time: 30 minutes
  Computed values:
    Merit ceiling: 32120
    Maximum decay: 12453
```

Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

CHAPTER 18

ES-IS Operational Commands

Copyright © 2018, Juniper Networks, Inc.

clear esis adjacency

Syntax

clear esis adjacency
<instance *instance-name*>
<interface *interface-name*)>
<*neighbor*>

Release Information Command introduced before Junos OS Release 7.4.

Description Clear End System-to-Intermediate System (ES-IS) adjacencies.

Options

none—Clear all ES-IS adjacencies.

instance *instance-name*—(Optional) Clear adjacencies for the specified routing instance only.

interface *interface-name*—(Optional) Clear adjacencies for the specified interface only.

neighbor—(Optional) Clear adjacencies for the specified neighbor only.

Required Privilege Level clear

Related Documentation

- show esis adjacency on page 1716

List of Sample Output clear esis adjacency on page 1714

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear esis adjacency

```
user@host> clear esis adjacency
```


Copyright © 2018, Juniper Networks, Inc.

clear esis statistics

Syntax

clear esis statistics
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 7.4.

Description

Clear End System-to-Intermediate System (ES-IS) packet statistics.

Options

none—Clear ES-IS packet statistics for all routing instances.

instance *instance-name*—(Optional) Clear ES-IS packet statistics for the specified routing instance only.

Required Privilege Level

clear

Related Documentation

- show esis statistics on page 1721

List of Sample Output

clear esis statistics on page 1715

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear esis statistics

```
user@host> clear esis statistics
```

show esis adjacency

Syntax

show esis adjacency
<brief | detail | extensive>
<*esis-neighbor-id*>
<instance *instance-name*>
<interface *interface-name*)>

Release Information Command introduced before Junos OS Release 7.4.

Description Display End System-to-Intermediate System (ES-IS) adjacencies.

Options

none—(Same as **brief**) Display all ES-IS adjacencies.

brief | detail | extensive—(Optional) Display the specified level of output.

esis-neighbor-id—(Optional) Display adjacencies for the specified neighbor's network service access point (NSAP) only.

instance *instance-name*—(Optional) Display adjacencies for the specified routing instance only.

interface *interface-name*—(Optional) Display adjacencies for the specified interface only.

Required Privilege Level view

Related Documentation

- clear esis adjacency on page 1714

List of Sample Output

show esis adjacency on page 1717
show esis adjacency brief on page 1717
show esis adjacency detail on page 1717
show esis adjacency extensive on page 1717

Output Fields Table 95 on page 1716 describes the output fields for the **show esis adjacency** command. Output fields are listed in the approximate order in which they appear.

Table 95: show esis adjacency Output Fields

Field Name	Field Description	Level of Output
Nbr Type	Type of network service access point (NSAP) of this neighbor.	**brief** none
NSAP/NET	NSAP of this neighbor.	All levels
Type	Type of NSAP of this neighbor.	**detail extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 95: show esis adjacency Output Fields (continued)

Field Name	Field Description	Level of Output
Hold (secs)	Holdtime interval advertised by this neighbor.	**brief** none
Interface	Interface through which the neighbor is reachable.	All levels
Advertised holdtime	Holdtime interval advertised by this neighbor.	**detail extensive**
Expires in	How long until the adjacency expires, in seconds.	**detail extensive**
SNPA	Subnetwork point of attachment (MAC address of the neighbor).	**detail extensive**
Transition log	List of recent transitions. • **When**—Time of advertisement from this neighbor. • **State**—State of the adjacency: **Up**, **Down**, **New**, **One-way**, **Initializing**, or **Rejected**. • **Event**—Event causing the state. • **Down reason**—Reason the adjacency is down.	**extensive**

Sample Output

show esis adjacency

```
user@host> show esis adjacency
Nbr  NSAP/NET                                          Hold  Interface
Type                                                  (secs)
IS   47.0005.80ff.f800.0000.0108.0001.0102.5501.6008    135 fe-0/0/0.0
```

show esis adjacency brief

The output for the **show esis adjacency brief** command is identical to that for the **show esis adjacency** command. For sample output, see show esis adjacency on page 1717.

show esis adjacency detail

```
user@host> show esis adjacency detail
NSAP/NET: 47.0005.80ff.f800.0000.0108.0001.0102.5501.6008, Type: IS
  Interface: fe-0/0/0.0, Advertised hold time: 180 secs, Expires in: 173 secs
  SNPA: 0:5:85:c1:73:71
```

show esis adjacency extensive

```
user@host> show esis adjacency extensive
NSAP/NET: 47.0005.80ff.f800.0000.0108.0001.0102.5501.6008, Type: IS
  Interface: fe-0/0/0.0, Advertised hold time: 180 secs, Expires in: 167 secs
  SNPA: 0:5:85:c1:73:71
  Transition log:
  When                  State        Event           Down reason
  Sun Nov 26 22:07:35   Up           Received ISH
```

Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

show esis interface

Syntax

show esis interface
<brief | detail | extensive>
<instance *instance-name*>
<interface *interface-name*)>

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display End System-to-Intermediate System (ES-IS) interface information.

Options

none—(Same as **brief**) Display information for all configured ES-IS interfaces.

brief | detail | extensive—(Optional) Display the specified level of output.

instance ***instance-name***—(Optional) Display configured interfaces for the specified routing instance only.

interface ***interface-name***—(Optional) Display information about the specified interface only.

Required Privilege Level

view

List of Sample Output

show esis interface on page 1720
show esis interface brief on page 1720
show esis interface detail on page 1720
show esis interface extensive on page 1720

Output Fields

Table 96 on page 1719 describes the output fields for the **show esis interface** command. Output fields are listed in the approximate order in which they appear.

Table 96: show esis interface Output Fields

Field Name	Field Description	Level of Output
Interface	Interface through which the adjacency is made.	All levels
Receives	Types of hello messages that are received.	All levels
Sends	Types of hello messages that are sent.	All levels
Hello interval	Interface's hello interval, in seconds.	All levels
Adjacencies or **Num Adj**	Number of adjacencies established on this interface.	All levels
Holdtime	Interface's hold time, in seconds.	**detail extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 96: show esis interface Output Fields (continued)

Field Name	Field Description	Level of Output
State	Internal implementation information.	**detail extensive**
End system configuration timer	Time, in seconds, for the end system to configure itself for ES-IS.	**detail extensive**
Interface index	Index value.	**detail extensive**
NET used in hello	Network entity title used in hello messages.	**detail extensive**

Sample Output

show esis interface

```
user@host> show esis interface
Interface                Receives     Sends  Hello Interval  Num Adj
fe-0/0/0.0                    ISH       ISH           60.00        1
lo0.0                         ISH         -           60.00        0
```

show esis interface brief

The output for the **show esis interface brief** command is identical to that for the **show esis interface** command. For sample output, see show esis interface on page 1720.

show esis interface detail

```
user@host> show esis interface detail
Interface: fe-0/0/0.0
  Receives: ISH, Sends: ISH, Hello interval: 60.00
  Adjacencies: 1, Holdtime: 180, End system configuration timer: 180
  Interface index: 68, State: 0x2
  NET used in hello: 47.0005.80ff.f800.0000.0108.0001.0102.5501.6007

Interface: lo0.0
  Receives: ISH, Sends:  - , Hello interval: 60.00
  Adjacencies: 0, Holdtime: 180, End system configuration timer: 180
  Interface index: 64, State: 0x2
  NET used in hello: 47.0005.80ff.f800.0000.0108.0001.0102.5501.6007
```

show esis interface extensive

The output for the **show esis interface extensive** command is identical to that for the **show esis interface detail** command. For sample output, see show esis interface detail on page 1720.

Copyright © 2018, Juniper Networks, Inc.

show esis statistics

Syntax

show esis statistics
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display End System-to-Intermediate System (ES-IS) packet statistics.

Options

none—Display ES-IS packet statistics for all routing instances.

instance *instance-name*—(Optional) Display ES-IS statistics for the specified routing instance only.

Required Privilege Level

view

Related Documentation

- clear esis statistics on page 1715

List of Sample Output

show esis statistics on page 1722

Output Fields

Table 97 on page 1721 describes the output fields for the **show esis statistics** command. Output fields are listed in the approximate order in which they appear.

Table 97: show esis statistics Output Fields

Field Name	Field Description
PDU type	Protocol data unit type.
Received	Number of PDUs received since IS-IS started or since the statistics were set to zero.
Processed	Number of PDUs received less the number dropped.
Drops	Number of PDUs dropped.
Sent	Number of PDUs transmitted since IS-IS started or since the statistics were set to zero.
Total packets received/sent	Total number of PDUs received and transmitted since IS-IS started or since the statistics were set to zero.

Sample Output

show esis statistics

```
user@host> show esis statistics
PDU type  Received Processed Drops    Sent
ESH              3         3     0       8
ISH             11        10     1       4
RD               0         0     0       0
Unknown          0         0     0       0
Totals          14        13     1      12
Total packets received: 14 sent: 0
```


Copyright © 2018, Juniper Networks, Inc.

CHAPTER 19

IP Multicast Operational Commands

- show dvmrp prefix
- show dvmrp prunes
- show igmp group
- show igmp interface
- show igmp snooping interface
- show igmp snooping membership
- show igmp snooping statistics
- show igmp statistics
- show mld group
- show mld interface
- show mld statistics
- show msdp
- show msdp source
- show msdp source-active
- show msdp statistics
- show multicast backup-pe-groups
- show multicast flow-map
- show multicast forwarding-cache statistics
- show multicast interface
- show multicast mrinfo
- show multicast next-hops
- show multicast pim-to-igmp-proxy
- show multicast pim-to-mld-proxy
- show multicast route
- show multicast rpf
- show multicast scope
- show multicast snooping next-hops
- show multicast sessions
- show multicast snooping route
- show multicast statistics
- show multicast usage
- show pim bootstrap
- show pim interfaces
- show pim join
- show pim mdt
- show pim mdt data-mdt-joins
- show pim mdt data-mdt-limit

Copyright © 2018, Juniper Networks, Inc.

- show pim neighbors
- show pim rps
- show pim snooping interfaces
- show pim snooping join
- show pim snooping neighbors
- show pim snooping statistics
- show pim source
- show pim statistics
- show sap listen
- test msdp

Copyright © 2018, Juniper Networks, Inc.

clear amt statistics

Syntax

clear amt statistics
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Release Information Command introduced in JUNOS Release 10.2.

Description Clear Automatic Multicast Tunneling (AMT) statistics.

Options

none—Clear the multicast statistics for all AMT tunnel interfaces.

instance *instance-name*—(Optional) Clear AMT multicast statistics for the specified instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level clear

Related Documentation

- show amt statistics on page 1764

List of Sample Output clear amt statistics on page 1726

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear amt statistics

```
user@host> clear amt statistics
```

 Copyright © 2018, Juniper Networks, Inc.

clear amt tunnel

Syntax

clear amt tunnel
<gateway *gateway-ip-addr*> <port *port-number*>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<statistics>
<tunnel-interface *interface-name*>

Release Information

Command introduced in JUNOS Release 10.2.

Description

Clear the Automatic Multicast Tunneling (AMT) multicast state. Optionally, clear AMT protocol statistics.

Options

none—Clear multicast state for all AMT tunnel interfaces.

gateway *gateway-ip-addr* port *port-number*—(Optional) Clear the AMT multicast state for the specified gateway address. If no port is specified, clear the AMT multicast state for all AMT gateways with the given IP address.

instance *instance-name*—(Optional) Clear the AMT multicast state for the specified instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

statistics—(Optional) Clear multicast statistics for all AMT tunnels or for specified tunnels.

tunnel-interface *interface-name*—(Optional) Clear the AMT multicast state for the specified AMT tunnel interface.

Required Privilege Level

clear

Related Documentation

- show amt tunnel on page 1769

List of Sample Output

clear amt tunnel on page 1727
clear amt tunnel statistics gateway-address on page 1728

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear amt tunnel

```
user@host> clear amt tunnel
```

Copyright © 2018, Juniper Networks, Inc.

clear amt tunnel statistics gateway-address

```
user@host> clear amt tunnel statistics gateway-address 100.31.1.21 port 4000
```


Copyright © 2018, Juniper Networks, Inc.

clear igmp membership

List of Syntax

Syntax on page 1729
Syntax (EX Series Switch and the QFX Series) on page 1729

Syntax

clear igmp membership
<all>
<group *address-range*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)

clear igmp membership
<group *address-range*>
<interface *interface-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Clear Internet Group Management Protocol (IGMP) group members.

Options

all—Clear IGMP members for groups and interfaces in the master instance.

group ***address-range***—(Optional) Clear all IGMP members that are in a particular address range. An example of a range is **233.252/16**. If you omit the destination prefix length, the default is **/32**.

interface ***interface-name***—(Optional) Clear all IGMP group members on an interface.

logical-system (all | ***logical-system-name*****)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Related Documentation

- show igmp group on page 1783
- show igmp interface on page 1787

List of Sample Output

clear igmp membership all on page 1730
clear igmp membership interface on page 1730
clear igmp membership group on page 1731

Output Fields

See show igmp group for an explanation of output fields.

Sample Output

clear igmp membership all

The following sample output displays IGMP group information before and after the **clear igmp membership** command is entered:

```
user@host> show igmp group
Interface         Group            Last Reported   Timeout
so-0/0/0          198.51.100.253    203.0.113.1          186
so-0/0/0          198.51.100.254    203.0.113.1          186
so-0/0/0          198.51.100.255    203.0.113.1          187
so-0/0/0          198.51.100.240    203.0.113.1          188
local             198.51.100.6      (null)                 0
local             198.51.100.5      (null)                 0
local             198.51.100.25     (null)                 0
local             198.51.100.22     (null)                 0
local             198.51.100.2      (null)                 0
local             198.51.100.13     (null)                 0
```

```
user@host> clear igmp membership all
Clearing Group Membership Info for so-0/0/0
Clearing Group Membership Info for so-1/0/0
Clearing Group Membership Info for so-2/0/0
```

```
user@host> show igmp group
Interface         Group            Last Reported   Timeout
local             198.51.100.6              (null)                0
local             198.51.100.5              (null)                0
local             198.51.100.254            (null)                0
local             198.51.100.255            (null)                0
local             198.51.100.2              (null)                0
local             198.51.100.13             (null)                0
```

clear igmp membership interface

The following sample output displays IGMP group information before and after the **clear igmp membership interface** command is issued:

```
user@host> show igmp group
Interface         Group            Last Reported   Timeout
so-0/0/0          198.51.100.253    203.0.113.1          210
so-0/0/0          198.51.100.200    203.0.113.1          210
so-0/0/0          198.51.100.255    203.0.113.1          215
so-0/0/0          198.51.100.254    203.0.113.1          216
local             198.51.100.6      (null)                 0
local             198.51.100.5      (null)                 0
local             198.51.100.254    (null)                 0
local             198.51.100.255    (null)                 0
local             198.51.100.2      (null)                 0
local             198.51.100.13     (null)                 0
```

```
user@host> clear igmp membership interface so-0/0/0
Clearing Group Membership Info for so-0/0/0
```

```
user@host> show igmp group
```

 Copyright © 2018, Juniper Networks, Inc.

```
Interface          Group            Last Reported   Timeout
local              198.51.100.6        (null)              0
local              198.51.100.5        (null)              0
local              198.51.100.254      (null)              0
local              198.51.100.255      (null)              0
local              198.51.100.2        (null)              0
local              198.51.100.13       (null)              0
```

clear igmp membership group

The following sample output displays IGMP group information before and after the **clear igmp membership group** command is entered:

```
user@host> show igmp group
Interface          Group            Last Reported   Timeout
so-0/0/0           198.51.100.253   203.0.113.1          210
so-0/0/0           198.51.100.25    203.0.113.1          210
so-0/0/0           198.51.100.255   203.0.113.1          215
so-0/0/0           198.51.100.254   203.0.113.1          216
local              198.51.100.6     (null)                 0
local              198.51.100.5     (null)                 0
local              198.51.100.254   (null)                 0
local              198.51.100.25    (null)                 0
local              198.51.100.2     (null)                 0
local              198.51.100.13    (null)                 0

user@host> clear igmp membership group 233.252/16
Clearing Group Membership Range 198.51.100.0/16 on so-0/0/0
Clearing Group Membership Range 198.51.100.0/16 on so-1/0/0
Clearing Group Membership Range 198.51.100.0/16 on so-2/0/0

user@host> show igmp group
Interface          Group            Last Reported   Timeout
so-0/0/0           198.51.100.255   203.0.113.1          231
so-0/0/0           198.51.100.254   203.0.113.1          233
so-0/0/0           198.51.100.253   203.0.113.1          236
local              198.51.100.6     (null)                 0
local              198.51.100.5     (null)                 0
local              198.51.100.254   (null)                 0
local              198.51.100.255   (null)                 0
local              198.51.100.2     (null)                 0
local              198.51.100.13    (null)                 0
```

Copyright © 2018, Juniper Networks, Inc.

clear igmp snooping membership

Syntax

clear igmp snooping membership
<vlan *vlan-name*>
<group | source *address*>
<instance *instance-name*>
<interface *interface-name*>
<learning-domain *learning-domain-name*>
<logical-system *logical-system-name*>
<vlan-id *vlan-identifier*>

Release Information

Command introduced in Junos OS Release 8.5.
Command introduced in Junos OS Release 18.1R1 for the SRX1500 devices.

Description

Clear IGMP snooping dynamic membership information from the multicast forwarding table.

Options

none—Clear IGMP snooping membership for all supported address families on all interfaces.

vlan *vlan-name* —(Optional) Clear dynamic membership information for the specified VLAN.

group | source *address*—(Optional) Clear IGMP snooping membership for the specified multicast group or source address.

instance *instance-name*—(Optional) Clear IGMP snooping membership for the specified instance.

interface *interface-name*—(Optional) Clear IGMP snooping membership on a specific interface.

learning-domain *learning-domain-name*—(Optional) Perform this operation on all learning domains or on a particular learning domain.

logical-system *logical-system-name*—(Optional) Display information about a particular logical system, or for all logical systems.

vlan-id *vlan-identifier*—(Optional) Perform this operation on a particular VLAN.

Required Privilege Level

clear

Related Documentation

• show igmp snooping membership on page 1796

List of Sample Output

clear igmp snooping membership on page 1733

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

clear igmp snooping membership

user@host> **clear igmp snooping membership**

clear igmp snooping statistics

Syntax

clear igmp snooping statistics
<instance *instance-name*>
<interface *interface-name*>
<learning-domain (all | *learning-domain-name*)>
<logical-system *logical-system-name*>

Release Information

Command introduced in Junos OS Release 8.5.
Command introduced in Junos OS Release 18.1R1 for the SRX1500 devices.

Description

Clear IP IGMP snooping statistics.

Options

none—Clear IGMP snooping statistics for all supported address families on all interfaces.

instance *instance-name*—(Optional) Clear IGMP snooping statistics for the specified instance.

interface *interface-name*—(Optional) Clear IGMP snooping statistics on a specific interface.

learning-domain (all | *learning-domain-name*)—(Optional) Perform this operation on all learning domains or on a particular learning domain.

logical-system *logical-system-name*—(Optional) Delete the IGMP snooping statistics for a given logical system or for all logical systems.

Required Privilege Level

clear

Related Documentation

- show igmp snooping statistics on page 1801

List of Sample Output

clear igmp snooping statistics on page 1734

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear igmp snooping statistics

```
user@host> clear igmp snooping statistics
```


Copyright © 2018, Juniper Networks, Inc.

clear igmp statistics

List of Syntax

Syntax on page 1735
Syntax (EX Series Switches) on page 1735

Syntax

clear igmp statistics
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

clear igmp statistics
<interface *interface-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Clear Internet Group Management Protocol (IGMP) statistics.

Options

none—Clear IGMP statistics on all interfaces.

interface *interface-name*—(Optional) Clear IGMP statistics for the specified interface only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Related Documentation

- show igmp statistics on page 1806

List of Sample Output

clear igmp statistics on page 1735

Output Fields

See **show igmp statistics** for an explanation of output fields.

Sample Output

clear igmp statistics

The following sample output displays IGMP statistics information before and after the **clear igmp statistics** command is entered:

```
user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type       Received        Sent  Rx errors
Membership Query            8883         459          0
V1 Membership Report           0           0          0
```

```
DVMRP                       19784      35476          0
PIM V1                      18310          0          0
Cisco Trace                     0          0          0
V2 Membership Report            0          0          0
Group Leave                     0          0          0
Mtrace Response                 0          0          0
Mtrace Request                  0          0          0
Domain Wide Report              0          0          0
V3 Membership Report            0          0          0
Other Unknown types                                   0
IGMP v3 unsupported type                              0
IGMP v3 source required for SSM                       0
IGMP v3 mode not applicable for SSM                   0

IGMP Global Statistics
Bad Length                      0
Bad Checksum                    0
Bad Receive If                  0
Rx non-local                 1227
```

```
user@host> clear igmp statistics
user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received       Sent  Rx errors
Membership Query                0          0          0
V1 Membership Report            0          0          0
DVMRP                           0          0          0
PIM V1                          0          0          0
Cisco Trace                     0          0          0
V2 Membership Report            0          0          0
Group Leave                     0          0          0
Mtrace Response                 0          0          0
Mtrace Request                  0          0          0
Domain Wide Report              0          0          0
V3 Membership Report            0          0          0
Other Unknown types                                   0
IGMP v3 unsupported type                              0
IGMP v3 source required for SSM                       0
IGMP v3 mode not applicable for SSM                   0
IGMP Global Statistics
Bad Length                      0
Bad Checksum                    0
Bad Receive If                  0
Rx non-local                    0
```

 Copyright © 2018, Juniper Networks, Inc.

clear mld membership

Syntax

clear mld membership
<all>
<group *group-name*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Release Information Command introduced before Junos OS Release 7.4.

Description Clear Multicast Listener Discovery (MLD) group membership.

Options

all—Clear MLD memberships for groups and interfaces in the master instance.

group *group-name*—(Optional) Clear MLD membership for the specified group.

interface *interface-name*—(Optional) Clear MLD group membership for the specified interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level view

Related Documentation

- show mld group on page 1809

List of Sample Output clear mld membership all on page 1737

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear mld membership all

```
user@host> clear mld membership all
```

clear mld statistics

Syntax

clear mld statistics
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced before Junos OS Release 7.4.

Description

Clear Multicast Listener Discovery (MLD) statistics.

Options

none—(Same as **logical-system all**) Clear MLD statistics for all interfaces.

interface *interface-name*—(Optional) Clear MLD statistics for the specified interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Related Documentation

- show mld statistics on page 1817

List of Sample Output

clear mld statistics on page 1738

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear mld statistics

`user@host>` **clear mld statistics**

 Copyright © 2018, Juniper Networks, Inc.

clear msdp cache

Syntax

clear msdp cache
<all>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<peer *peer-address*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Clear the entries in the Multicast Source Discovery Protocol (MSDP) source-active cache.

Options

all— Clear all MSDP source-active cache entries in the master instance.

instance *instance-name*—(Optional) Clear entries for a specific MSDP instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

peer *peer-address*—(Optional) Clear the MSDP source-active cache entries learned from a specific peer.

Required Privilege Level

clear

Related Documentation

- show msdp source-active on page 1825

List of Sample Output

clear msdp cache all on page 1739

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear msdp cache all

```
user@host> clear msdp cache all
```

clear msdp statistics

Syntax

clear msdp statistics
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<peer *peer-address*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Clear Multicast Source Discovery Protocol (MSDP) peer statistics.

Options

none—Clear MSDP statistics for all peers.

instance *instance-name*—(Optional) Clear statistics for the specified instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

peer *peer-address*—(Optional) Clear the statistics for the specified peer.

Required Privilege Level

clear

Related Documentation

- show msdp statistics on page 1828

List of Sample Output

clear msdp statistics on page 1740

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear msdp statistics

`user@host>` **clear msdp statistics**

Copyright © 2018, Juniper Networks, Inc.

clear multicast bandwidth-admission

Syntax

clear multicast bandwidth-admission
<group *group-address*>
<inet | inet6>
<instance *instance-name*>
<interface *interface-name*>
<source *source-address*>

Release Information

Command introduced in Junos OS Release 8.3.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Reapply IP multicast bandwidth admissions.

Options

none—Reapply multicast bandwidth admissions for all IPv4 forwarding entries in the master routing instance.

group *group-address*—(Optional) Reapply multicast bandwidth admissions for the specified group.

inet—(Optional) Reapply multicast bandwidth admission settings for IPv4 flows.

inet6—(Optional) Reapply multicast bandwidth admission settings for IPv6 flows.

instance *instance-name*—(Optional) Reapply multicast bandwidth admission settings for the specified instance. If you do not specify an instance, the command applies to the master routing instance.

interface *interface-name*—(Optional) Examines the corresponding outbound interface in the relevant entries and acts as follows:

- If the interface is congested, and it was admitted previously, it is removed.
- If the interface was rejected previously, the **clear multicast bandwidth-admission** command enables the interface to be admitted as long as enough bandwidth exists on the interface.
- If you do not specify an interface, issuing the **clear multicast bandwidth-admission** command readmits any previously rejected interface for the relevant entries as long as enough bandwidth exists on the interface.

To manually reject previously admitted outbound interfaces, you must specify the interface.

source *source-address*—(Optional) Use with the **group** option to reapply multicast bandwidth admission settings for the specified (source, group) entry.

Copyright © 2018, Juniper Networks, Inc.

Required Privilege Level
clear

Related Documentation
- show multicast interface on page 1838

List of Sample Output
clear multicast bandwidth-admission on page 1742

Output Fields
When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear multicast bandwidth-admission

```
user@host> clear multicast bandwidth-admission
```


Copyright © 2018, Juniper Networks, Inc.

clear multicast forwarding-cache

Syntax

clear multicast forwarding-cache
<all>
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced in Junos OS Release 12.2.

Description

Clear IP multicast forwarding cache entries.

This command is not supported for next-generation multiprotocol BGP multicast VPNs (MVPNs).

Options

all—Clear all multicast forwarding cache entries in the master instance.

inet—(Optional) Clear multicast forwarding cache entries for IPv4 family addresses.

inet6—(Optional) Clear multicast forwarding cache entries for IPv6 family addresses.

instance *instance-name*—(Optional) Clear multicast forwarding cache entries on a specific routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Related Documentation

- show multicast forwarding-cache statistics on page 1836

List of Sample Output

clear multicast forwarding-cache all on page 1743

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear multicast forwarding-cache all

`user@host>` **clear multicast forwarding-cache all**

clear multicast scope

List of Syntax
Syntax on page 1744
Syntax (EX Series Switch and the QFX Series) on page 1744

Syntax

clear multicast scope
<inet | inet6>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)

clear multicast scope
<inet | inet6>
<interface *interface-name*>

Release Information

Command introduced in Junos OS Release 7.6.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 option introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Clear IP multicast scope statistics.

Options

none—(Same as **logical-system all**) Clear multicast scope statistics.

inet—(Optional) Clear multicast scope statistics for IPv4 family addresses.

inet6—(Optional) Clear multicast scope statistics for IPv6 family addresses.

interface *interface-name*—(Optional) Clear multicast scope statistics on a specific interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Related Documentation

- show multicast scope on page 1866

List of Sample Output

clear multicast scope on page 1745

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

clear multicast scope

```
user@host> clear multicast scope
```

clear multicast sessions

List of Syntax

Syntax on page 1746
Syntax (EX Series Switch and the QFX Series) on page 1746

Syntax

clear multicast sessions
<logical-system (all | *logical-system-name*)>
<*regular-expression*>

Syntax (EX Series Switch and the QFX Series)

clear multicast sessions
<*regular-expression*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Clear IP multicast sessions.

Options

none—(Same as **logical-system all**) Clear multicast sessions.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

regular-expression—(Optional) Clear only multicast sessions that contain the specified regular expression.

Required Privilege Level

clear

Related Documentation

- show multicast sessions on page 1871

List of Sample Output

clear multicast sessions on page 1746

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear multicast sessions

```
user@host> clear multicast sessions
```


Copyright © 2018, Juniper Networks, Inc.

clear multicast snooping statistics

Syntax

clear multicast snooping statistics
<instance *instance-name*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Release Information Command introduced in Junos OS Release 8.5.

Description Clear IP multicast snooping statistics.

Options

none—Clear multicast snooping statistics for all supported address families on all interfaces.

instance ***instance-name***—(Optional) Clear multicast snooping statistics for the specified instance.

interface ***interface-name***—(Optional) Clear multicast snooping statistics on a specific interface.

logical-system (all | ***logical-system-name*****)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level clear

List of Sample Output clear multicast snooping statistics on page 1747

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear multicast snooping statistics

`user@host>` **clear multicast snooping statistics**

Copyright © 2018, Juniper Networks, Inc.

clear multicast statistics

List of Syntax Syntax on page 1748
Syntax (EX Series Switch and the QFX Series) on page 1748

Syntax

clear multicast statistics
<inet | inet6>
<instance *instance-name*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)

clear multicast statistics
<inet | inet6>
<instance *instance-name*>
<interface *interface-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Clear IP multicast statistics.

Options

none—Clear multicast statistics for all supported address families on all interfaces.

inet—(Optional) Clear multicast statistics for IPv4 family addresses.

inet6—(Optional) Clear multicast statistics for IPv6 family addresses.

instance *instance-name*—(Optional) Clear multicast statistics for the specified instance.

interface *interface-name*—(Optional) Clear multicast statistics on a specific interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level clear

Related Documentation

- show multicast statistics on page 1879

List of Sample Output clear multicast statistics on page 1749

Output Fields When you enter this command, you are provided feedback on the status of your request.

 Copyright © 2018, Juniper Networks, Inc.

Sample Output

clear multicast statistics

```
user@host> clear multicast statistics
```

clear pim join

List of Syntax

Syntax on page 1750
Syntax (EX Series Switch and the QFX Series) on page 1750

Syntax

clear pim join
<all>
<*group-address*>
<bidirectional | dense | sparse>
<exact>
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<rp *ip-address/prefix* | source *ip-address/prefix*>
<sg | star-g>

Syntax (EX Series Switch and the QFX Series)

clear pim join
<all>
<*group-address*>
<dense | sparse>
<exact>
<inet | inet6>
<instance *instance-name*>
<rp *ip-address/prefix* | source *ip-address/prefix*>
<sg | star-g>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Multiple new filter options introduced in Junos OS Release 13.2.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Clear the Protocol Independent Multicast (PIM) join and prune states.

Options

all—To clear PIM join and prune states for all groups and family addresses in the master instance, you must specify “all”..

group-address—(Optional) Clear the PIM join and prune states for a group address.

bidirectional | dense | sparse—(Optional) Clear PIM bidirectional mode, dense mode, or sparse and source-specific multicast (SSM) mode entries.

exact—(Optional) Clear only the group that exactly matches the specified group address.

inet | inet6—(Optional) Clear the PIM entries for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Clear the entries for a specific PIM-enabled routing instance.

Copyright © 2018, Juniper Networks, Inc.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

rp *ip-address/prefix* | source *ip-address/prefix*—(Optional) Clear the PIM entries with a specified rendezvous point (RP) address and prefix or with a specified source address and prefix. You can omit the prefix.

sg | star-g—(Optional) Clear PIM (S,G) or (*,G) entries.

Additional Information The **clear pim join** command cannot be used to clear the PIM join and prune state on a backup Routing Engine when nonstop active routing is enabled.

Required Privilege Level clear

Related Documentation

- show pim join on page 1892

List of Sample Output

clear pim join all on page 1751
clear pim join inet6 all on page 1751
clear pim join inet6 star-g all on page 1751

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear pim join all

```
user@host> clear pim join all
Cleared 8 Join/Prune states
```

clear pim join inet6 all

```
user@host> clear pim join inet6 all
Cleared 4 Join/Prune states
```

clear pim join inet6 star-g all

```
user@host> clear pim join inet6 star-g all
Cleared 1 Join/Prune states
```

Copyright © 2018, Juniper Networks, Inc.

clear pim join-distribution

Syntax

clear pim join-distribution
<all>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced in Junos OS Release 10.0.

Description

Clear the PIM **join-redistribute** states.

Use the **show pim source** command to find out if there are multiple paths available for a source (for example, an RP).

When you include the **join-load-balance** statement in the configuration, the PIM join states are distributed evenly on available equal-cost multipath links. When an upstream neighbor link fails, Junos OS redistributes the PIM join states to the remaining links. However, when new links are added or the failed link is restored, the existing PIM joins are not redistributed to the new link. New flows will be distributed to the new links. However, in a network without new joins and prunes, the new link is not used for multicast traffic. The **clear pim join-distribution** command redistributes the existing flows to the new upstream neighbors. Redistributing the existing flows causes traffic to be disrupted, so we recommend that you run the **clear pim join-distribution** command during a maintenance window.

Options

all— (Optional) Clear the PIM **join-redistribute** states for all groups and family addresses in the master instance.

none— Automatically clear all PIM join/prune states.

instance *instance-name*—(Optional) Redistribute the join states for a specific PIM-enabled routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information

The **clear pim join-distribution** command cannot be used to redistribute the PIM join states on a backup Routing Engine when nonstop active routing is enabled.

Required Privilege Level

clear

Related Documentation

- show pim neighbors on page 1915
- show pim join on page 1892
- *join-load-balance*

Copyright © 2018, Juniper Networks, Inc.

List of Sample Output clear pim join-distribution all on page 1753

Output Fields When you enter this command, you are provided no feedback on the status of your request. You can enter the **show pim join** command before and after distributing the join state to verify the operation.

Sample Output

clear pim join-distribution all

```
user@host> clear pim join-distribution all
```

clear pim register

List of Syntax

Syntax on page 1754
Syntax (EX Series Switch and the QFX Series) on page 1754
Syntax (PTX Series) on page 1754

Syntax

```
clear pim register
<all>
<inet | inet6>
<instance instance-name>
<interface interface-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and the QFX Series)

```
clear pim register
<inet | inet6>
<instance instance-name>
<interface interface-name>
```

Syntax (PTX Series)

```
clear pim register
<inet | inet6>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Release Information

Command introduced in Junos OS Release 7.6.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Clear Protocol Independent Multicast (PIM) register message counters.

Options

all—Required to clear the PIM register message counters for all groups and family addresses in the master instance.

inet | inet6—(Optional) Clear PIM register message counters for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Clear register message counters for a specific PIM-enabled routing instance.

interface *interface-name*—(Optional) Clear PIM register message counters for a specific interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information

The **clear pim register** command cannot be used to clear the PIM register state on a backup Routing Engine when nonstop active routing is enabled.

Copyright © 2018, Juniper Networks, Inc.

Required Privilege Level clear

Related Documentation

- show pim statistics on page 1947

List of Sample Output clear pim register all on page 1755

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear pim register all

```
user@host> clear pim register all
```

clear pim snooping join

Syntax

clear pim snooping join
<instance *instance-name*>
<logical-system *logical-system-name*>
<vlan-id *vlan-id*>

Release Information

Command introduced in Junos OS Release 12.3 for MX Series 3D Universal Edge devices.
Command introduced in Junos OS Release 13.2 for M Series Multiservice Edge devices.

Description

Clear information about Protocol Independent Multicast (PIM) snooping joins.

Options

none—Display detailed information.

instance ***instance-name***—(Optional) Clear PIM snooping join information for the specified routing instance.

logical-system ***logical-system-name***—(Optional) Delete the IGMP snooping statistics for a given logical system or for all logical systems.

vlan-id ***vlan-identifier***—(Optional) Clear PIM snooping join information for the specified VLAN.

Required Privilege Level

view

Related Documentation

- *PIM Snooping for VPLS*

List of Sample Output

clear pim snooping join on page 1756

Output Fields

See show pim snooping join for an explanation of the output fields.

Sample Output

clear pim snooping join

The following sample output displays information about PIM snooping joins before and after the **clear pim snooping join** command is entered:

```
user@host> show pim snooping join extensive
Instance: vpls1
Learning-Domain: vlan-id 10
Learning-Domain: vlan-id 20

Group: 198.51.100.2
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 192.0.2.5, port: ge-1/3/7.20
Downstream port: ge-1/3/1.20
```

 Copyright © 2018, Juniper Networks, Inc.

```
Downstream neighbors:
192.0.2.2 State: Join Flags: SRW Timeout: 185

Group: 198.51.100.3
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 192.0.2.4, port: ge-1/3/5.20
Downstream port: ge-1/3/3.20
Downstream neighbors:
192.0.2.3 State: Join Flags: SRW Timeout: 175
```

```
user@host> clear pim snooping join
Clearing the Join/Prune state for 203.0.113.0/24
Clearing the Join/Prune state for 203.0.113.0/24
```

```
user@host> show pim snooping join extensive
Instance: vpls1
Learning-Domain: vlan-id 10
Learning-Domain: vlan-id 20
```

clear pim snooping statistics

Syntax

clear pim snooping statistics
<instance *instance-name*>
<interface *interface-name*>
<logical-system *logical-system-name*>
<vlan-id *vlan-id*>

Release Information

Command introduced in Junos OS Release 12.3 for MX Series 3D Universal Edge devices.
Command introduced in Junos OS Release 13.2 for M Series Multiservice Edge devices.

Description

Clear Protocol Independent Multicast (PIM) snooping statistics.

Options

none—Clear PIM snooping statistics for all family addresses, instances, and interfaces.

instance ***instance-name***—(Optional) Clear statistics for a specific PIM-snooping-enabled routing instance.

interface ***interface-name***—(Optional) Clear PIM snooping statistics for a specific interface.

logical-system ***logical-system-name***—(Optional) Delete the IGMP snooping statistics for a given logical system or for all logical systems.

vlan-id ***vlan-identifier***—(Optional) Clear PIM snooping statistics information for the specified VLAN.

Required Privilege Level

clear

Related Documentation

- *PIM Snooping for VPLS*

List of Sample Output

clear pim snooping statistics on page 1758

Output Fields

See show pim snooping statistics for an explanation of the output fields.

Sample Output

clear pim snooping statistics

The following sample output displays PIM snooping statistics before and after the **clear pim snooping statistics** command is entered:

```
user@host> show pim snooping statistics
Instance: vpls1
Learning-Domain: vlan-id 10

Tx J/P messages 0
RX J/P messages 660
Rx J/P messages -- seen 0
```

 Copyright © 2018, Juniper Networks, Inc.

```
Rx J/P messages -- received 660
Rx Hello messages 1396
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0

Learning-Domain: vlan-id 20
```

```
user@host> clear pim snooping statistics
user@host> show pim snooping statistics
Instance: vpls1
Learning-Domain: vlan-id 10

Tx J/P messages 0
RX J/P messages 0
Rx J/P messages -- seen 0
Rx J/P messages -- received 0
Rx Hello messages 0
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0

Learning-Domain: vlan-id 20
```

Copyright © 2018, Juniper Networks, Inc.

clear pim statistics

List of Syntax Syntax on page 1760
Syntax (EX Series Switch and the QFX Series) on page 1760

Syntax

```
clear pim statistics
<inet | inet6>
<instance instance-name>
<interface interface-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and the QFX Series)

```
clear pim statistics
<inet | inet6>
<instance instance-name>
<interface interface-name>
```

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Clear Protocol Independent Multicast (PIM) statistics.

Options

none—Clear PIM statistics for all family addresses, instances, and interfaces.

inet | inet6—(Optional) Clear PIM statistics for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Clear statistics for a specific PIM-enabled routing instance.

interface *interface-name*—(Optional) Clear PIM statistics for a specific interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information The **clear pim statistics** command cannot be used to clear the PIM statistics on a backup Routing Engine when nonstop active routing is enabled.

Required Privilege Level clear

Related Documentation

- show pim statistics on page 1947

List of Sample Output clear pim statistics on page 1761

Output Fields See show pim statistics for an explanation of output fields.

 Copyright © 2018, Juniper Networks, Inc.

Sample Output

clear pim statistics

The following sample output displays PIM statistics before and after the **clear pim statistics** command is entered:

```
user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type       Received       Sent  Rx errors
Hello                         0          0          0
Register                      0          0          0
Register Stop                 0          0          0
Join Prune                    0          0          0
Bootstrap                     0          0          0
Assert                        0          0          0
Graft                         0          0          0
Graft Ack                     0          0          0
Candidate RP                  0          0          0
V1 Query                   2111       4222          0
V1 Register                   0          0          0
V1 Register Stop              0          0          0
V1 Join Prune             14200      13115          0
V1 RP Reachability            0          0          0
V1 Assert                     0          0          0
V1 Graft                      0          0          0
V1 Graft Ack                  0          0          0
PIM statistics summary for all interfaces:
Unknown type                            0
V1 Unknown type                         0
Unknown Version                         0
Neighbor unknown                        0
Bad Length                              0
Bad Checksum                            0
Bad Receive If                          0
Rx Intf disabled                     2007
Rx V1 Require V2                        0
Rx Register not RP                      0
RP Filtered Source                      0
Unknown Reg Stop                        0
Rx Join/Prune no state               1040
Rx Graft/Graft Ack no state             0
...

user@host> clear pim statistics
user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type       Received       Sent  Rx errors
Hello                         0          0          0
Register                      0          0          0
Register Stop                 0          0          0
Join Prune                    0          0          0
Bootstrap                     0          0          0
Assert                        0          0          0
Graft                         0          0          0
Graft Ack                     0          0          0
Candidate RP                  0          0          0
V1 Query                      1          0          0
```

Copyright © 2018, Juniper Networks, Inc.

```
V1 Register                   0           0           0
...
```

request pim multicast-tunnel rebalance

List of Syntax

Syntax on page 1763
Syntax (EX Series Switches) on page 1763

Syntax

request pim multicast-tunnel rebalance
<instance *instance-name*>
<logical-system (all | *logical-system-name*>

Syntax (EX Series Switches)

request pim multicast-tunnel rebalance
<instance *instance-name*>

Release Information

Command introduced in Junos OS Release 10.2.
Command introduced in Junos OS Release 10.2 for EX Series switches.

Description

Rebalance the assignment of multicast tunnel encapsulation interfaces across available tunnel-capable PICs or across a configured list of tunnel-capable PICs. You can determine whether a rebalance is necessary by running the **show pim interfaces instance *instance-name*** command.

Options

none—Re-create and rebalance all tunnel interfaces for all routing instances.

instance *instance-name*—Re-create and rebalance all tunnel interfaces for a specific instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

maintenance

Related Documentation

- show pim interfaces on page 1889
- *Load Balancing Multicast Tunnel Interfaces Among Available PICs*

Output Fields

This command produces no output. To verify the operation of the command, run the **show pim interface instance *instance-name*** before and after running the **request pim multicast-tunnel rebalance** command.

show amt statistics

Syntax

show amt statistics
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced in JUNOS Release 10.2.

Description

Display information about the Automatic Multicast Tunneling (AMT) protocol tunnel statistics.

Options

none—Display summary information about all AMT Protocol tunnels.

instance *instance-name*—(Optional) Display information for the specified instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- clear amt statistics on page 1726
- show amt summary on page 1767
- show amt tunnel on page 1769

List of Sample Output

show amt statistics on page 1765

Output Fields

Table 98 on page 1764 describes the output fields for the **show amt statistics** command. Output fields are listed in the approximate order in which they appear.

Table 98: show amt statistics Output Fields

Field Name	Field Description
AMT receive message count	Summary of AMT statistics for messages received on all interfaces. • **AMT relay discovery**—Number of AMT relay discovery messages received. • **AMT membership request**—Number of AMT membership request messages received. • **AMT membership update**—Number of AMT membership update messages received.
AMT send message count	Summary of AMT statistics for messages sent on all interfaces. • **AMT relay advertisement**—Number of AMT relay advertisement messages sent. • **AMT membership query**—Number of AMT membership query messages sent.

Copyright © 2018, Juniper Networks, Inc.

Table 98: show amt statistics Output Fields (continued)

Field Name	Field Description
AMT error message count	Summary of AMT statistics for error messages received on all interfaces. • **AMT incomplete packet**—Number of messages received with length errors so severe that further classification could not occur. • **AMT invalid mac**—Number of messages received with an invalid message authentication code (MAC). • **AMT unexpected type**—Number of messages received with an unknown message type specified. • **AMT invalid relay discovery address**—Number of AMT relay discovery messages received with an address other than the configured anycast address. • **AMT invalid membership request address**—Number of AMT membership request messages received with an address other than the configured AMT local address. • **AMT invalid membership update address**—Number of AMT membership update messages received with an address other than the configured AMT local address. • **AMT incomplete relay discovery messages**—Number of AMT relay discovery messages received that are not fully formed. • **AMT incomplete membership request messages**—Number of AMT membership request messages received that are not fully formed. • **AMT incomplete membership update messages**—Number of AMT membership update messages received that are not fully formed. • **AMT no active gateway**—Number of AMT membership update messages received for a tunnel that does not exist for the gateway that sent the message. • **AMT invalid inner header checksum**—Number of AMT membership update messages received with an invalid IP checksum. • **AMT gateways timed out**—Number of gateways that timed out because of inactivity.

Sample Output

show amt statistics

```
user@host> show amt statistics

AMT receive message count
AMT relay advertisement                      :          2
AMT membership request                       :          5
AMT membership update                        :          5

AMT send message count
AMT relay advertisement                      :          2
AMT membership query                         :          5

AMT error message count
AMT incomplete packet                        :          0
AMT invalid mac                              :          0
AMT unexpected type                          :          0
AMT invalid relay discovery address          :          0
AMT invalid membership request address       :          0
AMT invalid membership update address        :          0
AMT incomplete relay discovery messages      :          0
AMT incomplete membership request messages   :          0
AMT incomplete membership update messages    :          0
AMT no active gateway                        :          0
```

Copyright © 2018, Juniper Networks, Inc.

```
AMT invalid inner header checksum          :          0
AMT gateways timed out                     :          0
```

 Copyright © 2018, Juniper Networks, Inc.

show amt summary

Syntax

show amt summary
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced in Junos OS Release 10.2.

Description

Display summary information about the Automatic Multicast Tunneling (AMT) protocol.

Options

none—Display summary information about all AMT protocol instances.

instance *instance-name*—(Optional) Display information for the specified instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- clear amt tunnel on page 1727
- show amt statistics on page 1764
- show amt tunnel on page 1769

List of Sample Output

show amt summary on page 1768

Output Fields

Table 99 on page 1767 describes the output fields for the **show amt summary** command. Output fields are listed in the approximate order in which they appear.

Table 99: show amt summary Output Fields

Field Name	Field Description	Level of Output
AMT anycast prefix	Prefix advertised by unicast routing protocols to route AMT discovery messages to the router from nearby AMT gateways.	All levels
AMT anycast address	Anycast address configured from which the anycast prefix is derived.	All levels
AMT local address	Local unique AMT relay IP address configured. Used to send AMT relay advertisement messages, it is the IP source address of AMT control messages and the source address of the data tunnel encapsulation.	All levels
AMT tunnel limit	Maximum number of AMT tunnels that can be created.	All levels
active tunnels	Number of active AMT tunnel interfaces.	All levels

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show amt summary

```
user@host> show amt summary
 AMT anycast prefix : 20.0.0.4/32
 AMT anycast address : 20.0.0.4
 AMT local address : 20.0.0.4
 AMT tunnel limit : 1000, active tunnels : 2
```

 Copyright © 2018, Juniper Networks, Inc.

show amt tunnel

Syntax

show amt tunnel
<brief | detail>
<gateway-address *gateway-ip-address*> <port *port-number*>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<tunnel-interface *interface-name*>

Release Information

Command introduced in Junos OS Release 10.2.

Description

Display information about the Automatic Multicast Tunneling (AMT) dynamic tunnels.

Options

none—Display summary information about all AMT protocol instances.

brief | detail—(Optional) Display the specified level of detail.

gateway-address *gateway-ip-address* port *port-number*—(Optional) Display information for the specified AMT gateway only. If no port is specified, display information for all AMT gateways with the given IP address.

instance *instance-name*—(Optional) Display information for the specified instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

tunnel-interface *interface-name*—(Optional) Display information for the specified AMT tunnel interface only.

Required Privilege Level

view

Related Documentation

- clear amt tunnel on page 1727
- show amt statistics on page 1764
- show amt summary on page 1767

List of Sample Output

show amt tunnel on page 1770
show amt tunnel detail on page 1771
show amt tunnel tunnel-interface on page 1771
show amt tunnel gateway-address on page 1771
show amt tunnel gateway-address detail on page 1771

Output Fields

Table 100 on page 1770 describes the output fields for the **show amt tunnel** command. Output fields are listed in the approximate order in which they appear.

Table 100: show amt tunnel Output Fields

Field Name	Field Description	Level of Output
AMT gateway address	Address of the AMT gateway that is being connected by the AMT tunnel.	All levels
port	Client port used by the AMT tunnel.	All levels
AMT tunnel interface	Dynamically created AMT logical interfaces used by the AMT tunnel in the format **ud-FPC/PIC/Port.unit**.	All levels
AMT tunnel state	State of the AMT tunnel. The state is normally **Active**. • **Active**—The tunnel is active. • **Pending**—The tunnel creation is pending. This is a transient state. • **Down**—The tunnel is in the down state. • **Graceful restart pending**—Graceful restart is in progress. • **Reviving**—The routing protocol daemon or Routing Engine was restarted (not gracefully). The tunnel remains in the reviving state until the AMT gateway sends a control message. When the message is received the tunnel is moved to the **Active** state. If no message is received before the AMT tunnel inactivity timer expires, the tunnel is deleted.	All levels
AMT tunnel inactivity timeout	Number of seconds since the most recent control message was received from an AMT gateway. If no message is received before the AMT tunnel inactivity timer expires, the tunnel is deleted.	All levels
Number of groups	Number of multicast groups using the tunnel.	All levels
Group	Multicast group address or addresses using the tunnel.	**detail**
Include Source	Multicast source address for each IGMPv3 group using the tunnel.	**detail**
AMT message count	Statistics for AMT messages: • **AMT Request**—Number of AMT relay tunnel request messages received. • **AMT membership update**—Number of AMT membership update messages received.	All levels

Sample Output

show amt tunnel

```
user@host> show amt tunnel
AMT gateway address : 11.11.11.2, port : 2268
 AMT tunnel interface : ud-5/1/10.1120256
 AMT tunnel state : Active
 AMT tunnel inactivity timeout : 15
 Number of groups : 1

AMT message count:
AMT Request    AMT membership update
2              2
```

 Copyright © 2018, Juniper Networks, Inc.

show amt tunnel detail

```
user@host> show amt tunnel detail
AMT gateway address : 11.11.11.2, port : 2268
 AMT tunnel interface : ud-5/3/10.1120512
 AMT tunnel state : Active
 AMT tunnel inactivity timeout : 62
 Number of groups : 1
  Group: 226.2.3.2

AMT message count:
AMT Request    AMT membership update
2              2

 AMT gateway address : 11.11.11.3, port : 2268
 AMT tunnel interface : ud-5/2/10.1120513
 AMT tunnel state : Active
 AMT tunnel inactivity timeout : 214
 Number of groups : 1
  Group: 226.2.3.3

AMT message count:
AMT Request    AMT membership update
2              2
```

show amt tunnel tunnel-interface

```
user@host> show amt tunnel tunnel-interface ud-5/3/10.1120512
 AMT gateway address : 11.11.11.2, port : 2268
 AMT tunnel interface : ud-5/3/10.1120512
 AMT tunnel state : Active
 AMT tunnel inactivity timeout : 145
 Number of groups : 1

AMT message count:
AMT Request    AMT membership update
2              2
```

show amt tunnel gateway-address

```
user@host> show amt tunnel gateway-address 11.11.11.3 port 2268
AMT gateway address : 11.11.11.3, port : 2268
 AMT tunnel interface : ud-5/2/10.1120513
 AMT tunnel state : Active
 AMT tunnel inactivity timeout : 214
 Number of groups : 1
  Group: 226.2.3.3

AMT message count:
AMT Request    AMT membership update
2              2
```

show amt tunnel gateway-address detail

```
user@host> show amt tunnel gateway-address 11.11.11.2 detail
AMT gateway address : 11.11.11.2, port : 2268
 AMT tunnel interface : ud-5/3/10.1120512
 AMT tunnel state : Active
 AMT tunnel inactivity timeout : 234
 Number of groups : 1
```

Copyright © 2018, Juniper Networks, Inc.

```
  Group: 226.2.3.2

AMT message count:
AMT Request    AMT membership update
2              2
```


Copyright © 2018, Juniper Networks, Inc.

show dvmrp interfaces

Syntax

show dvmrp interfaces
<logical-system (all | *logical-system-name*)>

Release Information

NOTE: Distance Vector Multicast Routing Protocol (DVMRP) was deprecated in Junos OS Release 16.1. Although DVMRP commands continue to be available and configurable in the CLI, they are no longer visible and are scheduled for removal in a subsequent release.

Command introduced before Junos OS Release 7.4.

Description

Display information about Distance Vector Multicast Routing Protocol (DVMRP)–enabled interfaces.

Options

none—(Same as **logical-system all**) Display information about DVMRP-enabled interfaces.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show dvmrp interfaces on page 1775

Output Fields

Table 101 on page 1773 describes the output fields for the **show dvmrp interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 101: show dvmrp interfaces Output Fields

Field Name	Field Description
Interface	Name of the interface.
State	State of the interface: **up** or **down**.
Leaf	Whether the interface is a leaf (that is, whether it has no neighbors) or whether it has neighbors.
Metric	Interface metric: a value from **1** through **31**.
Announce	Number of routes the interface is announcing.

Copyright © 2018, Juniper Networks, Inc.

Table 101: show dvmrp interfaces Output Fields (continued)

Field Name	Field Description
Mode	DVMRP mode: • **Forwarding**—DVMRP does both the routing and the multicast data forwarding. • **Unicast-routing**—DVMRP does only the routing. Forwarding of the multicast data packets can be done by enabling PIM on the interface.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show dvmrp interfaces

```
user@host> show dvmrp interfaces
Interface State Leaf Metric Announce Mode
fxp0.0    Up    N    1    4 Forwarding
fxp1.0    Up    N    1    4 Forwarding
fxp2.0    Up    N    1    3 Forwarding
lo0.0     Up    Y    1    0 Unicast-routing
```

show dvmrp neighbors

Syntax

show dvmrp neighbors
<logical-system (all | *logical-system-name*)>

Release Information

NOTE: Distance Vector Multicast Routing Protocol (DVMRP) was deprecated in Junos OS Release 16.1. Although DVMRP commands continue to be available and configurable in the CLI, they are no longer visible and are scheduled for removal in a subsequent release.

Command introduced before Junos OS Release 7.4.

Description

Display information about Distance Vector Multicast Routing Protocol (DVMRP) neighbors.

Options

none—(Same as **logical-system all**) Display information about DVMRP neighbors.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show dvmrp neighbors on page 1777

Output Fields

Table 102 on page 1776 describes the output fields for the **show dvmrp neighbors** command. Output fields are listed in the approximate order in which they appear.

Table 102: show dvmrp neighbors Output Fields

Field Name	Field Description
Neighbor	Address of the neighboring DVMRP router.
Interface	Interface through which the neighbor is reachable.
Version	Version of DVMRP that the neighbor is running, in the format ***majorminor***.
Flags	Information about the neighbor: • **1**—One way. The local router has seen the neighbor, but the neighbor has not seen the local router. • **G**—Neighbor supports generation ID. • **L**—Neighbor is a leaf router. • **M**—Neighbor supports mtrace. • **N**—Neighbor supports netmask in prune messages and graft messages. • **P**—Neighbor supports pruning. • **S**—Neighbor supports SNMP.

Copyright © 2018, Juniper Networks, Inc.

Table 102: show dvmrp neighbors Output Fields (continued)

Field Name	Field Description
Routes	Number of routes learned from the neighbor.
Timeout	How long until the DVMRP neighbor information times out, in seconds.
Transitions	Number of generation ID changes that have occurred since the local router learned about the neighbor.

Sample Output

show dvmrp neighbors

```
user@host> show dvmrp neighbors
Neighbor         Interface        Version Flags   Routes Timeout Transitions
192.168.1.1     ipip.0             3.255  PGM          3      28           1
```

show dvmrp prefix

Syntax

show dvmrp prefix
<brief | detail>
<logical-system (all | *logical-system-name*)>
<*prefix*>

Release Information

NOTE: Distance Vector Multicast Routing Protocol (DVMRP) was deprecated in Junos OS Release 16.1. Although DVMRP commands continue to be available and configurable in the CLI, they are no longer visible and are scheduled for removal in a subsequent release.

Command introduced before Junos OS Release 7.4.

Description

Display information about Distance Vector Multicast Routing Protocol (DVMRP) prefixes.

Options

none—Display standard information about all DVMRP prefixes.

brief | detail—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

prefix—(Optional) Display information about specific prefixes.

Required Privilege Level

view

List of Sample Output

show dvmrp prefix on page 1780
show dvmrp prefix brief on page 1780
show dvmrp prefix detail on page 1780

Output Fields

Table 103 on page 1778 describes the output fields for the **show dvmrp prefix** command. Output fields are listed in the approximate order in which they appear.

Table 103: show dvmrp prefix Output Fields

Field Name	Field Description	Level of Output
Prefix	DVMRP route.	All levels
Next hop	Next hop from which the route was learned.	All levels
Age	Last time that the route was refreshed.	All levels
multicast-group	Multicast group address.	**detail**

 Copyright © 2018, Juniper Networks, Inc.

Table 103: show dvmrp prefix Output Fields (continued)

Field Name	Field Description	Level of Output
Prunes sent	Number of prune messages sent to the multicast group.	**detail**
Grafts sent	Number of grafts sent to the multicast group.	**detail**
Cache lifetime	Lifetime of the group in the multicast cache, in seconds.	**detail**
Prune lifetime	Lifetime remaining and total lifetime of prune messages, in seconds.	**detail**

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show dvmrp prefix

```
user@host> show dvmrp prefix
Prefix                  Next hop          Age
10.38.0.0       /30 10.38.0.1         00:06:17
10.38.0.4       /30 10.38.0.5         00:06:13
10.38.0.8       /30 10.38.0.2         00:00:04
10.38.0.12      /30 10.38.0.6         00:00:04
10.255.14.114   /32 10.255.14.114     00:06:17
10.255.14.142   /32 10.38.0.2         00:00:04
10.255.14.144   /32 10.38.0.2         00:00:04
10.255.70.15    /32 10.38.0.6         00:00:04
192.168.14.0    /24 192.168.14.114    00:06:17
192.168.195.40 /30 192.168.195.41     00:06:17
192.168.195.92 /30 10.38.0.2          00:00:04
```

show dvmrp prefix brief

The output for the **show dvmrp prefix brief** command is identical to that for the **show dvmrp prefix** command.

show dvmrp prefix detail

```
user@host> show dvmrp prefix detail
Prefix                  Next hop          Age
10.38.0.0       /30 10.38.0.1         00:06:28
10.38.0.4       /30 10.38.0.5         00:06:24
10.38.0.8       /30 10.38.0.2         00:00:15
10.38.0.12      /30 10.38.0.6         00:00:15
10.255.14.114   /32 10.255.14.114     00:06:28
10.255.14.142   /32 10.38.0.2         00:00:15
10.255.14.144   /32 10.38.0.2         00:00:15
10.255.70.15    /32 10.38.0.6         00:00:15
192.168.14.0    /24 192.168.14.114    00:06:28
192.168.195.40 /30 192.168.195.41     00:06:28
192.168.195.92 /30 10.38.0.2          00:00:15
```


Copyright © 2018, Juniper Networks, Inc.

show dvmrp prunes

Syntax

show dvmrp prunes
<all | rx | tx>
<logical-system (all | *logical-system-name*)>

Release Information

NOTE: Distance Vector Multicast Routing Protocol (DVMRP) was deprecated in Junos OS Release 16.1. Although DVMRP commands continue to be available and configurable in the CLI, they are no longer visible and are scheduled for removal in a subsequent release.

Command introduced before Junos OS Release 7.4.

Description

Display information about active Distance Vector Multicast Routing Protocol (DVMRP) prune messages.

Options

none—Display received and transmitted DVMRP prune information.

all—(Optional) Display information about all received and transmitted prune messages.

rx—(Optional) Display information about received prune messages.

tx—(Optional) Display information about transmitted prune messages.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show dvmrp prunes on page 1782

Output Fields

Table 104 on page 1781 describes the output fields for the **show dvmrp prunes** command. Output fields are listed in the approximate order in which they appear.

Table 104: show dvmrp prunes Output Fields

Field Name	Field Description
Group	Group address.
Source prefix	Prefix for the prune.
Timeout	How long until the prune message expires, in seconds.
Neighbor	Neighbor to which the prune was sent or from which the prune was received.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show dvmrp prunes

```
user@host> show dvmrp prunes
Group            Source prefix      Timeout Neighbor
224.0.1.1        128.112.0.0    /12    7077 192.168.1.1
224.0.1.32       160.0.0.0      /3     7087 192.168.1.1
224.2.123.4      136.0.0.0      /5     6955 192.168.1.1
224.2.127.1      129.0.0.0      /8     7046 192.168.1.1
224.2.135.86     128.102.128.0  /17    7071 192.168.1.1
224.2.135.86     129.0.0.0      /8     7074 192.168.1.1
224.2.135.86     130.0.0.0      /7     7071 192.168.1.1
...
```

 Copyright © 2018, Juniper Networks, Inc.

show igmp group

List of Syntax

Syntax on page 1783
Syntax (EX Series Switch and the QFX Series) on page 1783

Syntax

show igmp group
<brief | detail>
<*group-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)

show igmp group
<brief | detail>
<*group-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display Internet Group Management Protocol (IGMP) group membership information.

Options

none—Display standard information about membership for all IGMP groups.

brief | detail—(Optional) Display the specified level of output.

group-name—(Optional) Display group membership for the specified IP address only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- clear igmp membership on page 1729

List of Sample Output

show igmp group (Include Mode) on page 1784
show igmp group (Exclude Mode) on page 1785
show igmp group brief on page 1785
show igmp group detail on page 1785

Output Fields

Table 105 on page 1784 describes the output fields for the **show igmp group** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2018, Juniper Networks, Inc.

Table 105: show igmp group Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface that received the IGMP membership report. A name of **local** indicates that the local routing device joined the group itself.	All levels
Group	Group address.	All levels
Group Mode	Mode the SSM group is operating in: **Include** or **Exclude**.	All levels
Source	Source address.	All levels
Source timeout	Time remaining until the group traffic is no longer forwarded. The timer is refreshed when a listener in include mode sends a report. A group in exclude mode or configured as a static group displays a zero timer.	**detail**
Last reported by	Address of the host that last reported membership in this group.	All levels
Timeout	Time remaining until the group membership is removed.	**brief none**
Group timeout	Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer.	**detail**
Type	Type of group membership: • **Dynamic**—Host reported the membership. • **Static**—Membership is configured.	All levels

Sample Output

show igmp group (Include Mode)

```
user@host> show igmp group
Interface: t1-0/1/0.0
    Group: 198.51.100.1
        Group mode: Include
        Source: 203.0.113.2
        Last reported by: 203.0.113.52
        Timeout:      24 Type: Dynamic
    Group: 198.51.100.1
        Group mode: Include
        Source: 203.0.113.3
        Last reported by: 203.0.113.52
        Timeout:      24 Type: Dynamic
    Group: 198.51.100.1
        Group mode: Include
        Source: 203.0.113.4
        Last reported by: 203.0.113.52
        Timeout:      24 Type: Dynamic
    Group: 198.51.100.2
        Group mode: Include
        Source: 203.0.113.4
        Last reported by: 203.0.113.52
```


Copyright © 2018, Juniper Networks, Inc.

```
        Timeout:      24 Type: Dynamic
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
    Group: 198.51.100.12
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: 198.51.100.22
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic
```

show igmp group (Exclude Mode)

```
user@host> show igmp group
Interface: t1-0/1/0.0
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
    Group: 198.51.100.2
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: 198.51.100.22
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic
```

show igmp group brief

The output for the **show igmp group brief** command is identical to that for the **show igmp group** command.

show igmp group detail

```
user@host> show igmp group detail
Interface: t1-0/1/0.0
    Group: 198.51.100.1
        Group mode: Include
        Source: 203.0.113.2
        Source timeout: 12
        Last reported by: 203.0.113.52
        Group timeout:       0 Type: Dynamic
    Group: 198.51.100.1
        Group mode: Include
        Source: 203.0.113.3
        Source timeout: 12
        Last reported by: 203.0.113.52
        Group timeout:       0 Type: Dynamic
    Group: 198.51.100.1
        Group mode: Include
        Source: 203.0.113.4
        Source timeout: 12
        Last reported by: 203.0.113.52
        Group timeout:       0 Type: Dynamic
    Group: 198.51.100.2
```

Copyright © 2018, Juniper Networks, Inc.

```
        Group mode: Include
        Source: 203.0.113.4
        Source timeout: 12
        Last reported by: 203.0.113.52
        Group timeout:       0 Type: Dynamic
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
    Group: 198.51.100.12
        Group mode: Exclude
        Source: 0.0.0.0
        Source timeout: 0
        Last reported by: Local
        Group timeout:       0 Type: Dynamic
    Group: 198.51.100.22
        Group mode: Exclude
        Source: 0.0.0.0
        Source timeout: 0
        Last reported by: Local
        Group timeout:       0 Type: Dynamic
```

 Copyright © 2018, Juniper Networks, Inc.

show igmp interface

List of Syntax

Syntax on page 1787
Syntax (EX Series Switches and the QFX Series) on page 1787

Syntax

show igmp interface
<brief | detail>
<*interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches and the QFX Series)

show igmp interface
<brief | detail>
<*interface-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display information about Internet Group Management Protocol (IGMP)-enabled interfaces.

Options

none—Display standard information about all IGMP-enabled interfaces.

brief | detail—(Optional) Display the specified level of output.

interface-name—(Optional) Display information about the specified IGMP-enabled interface only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- clear igmp membership on page 1729

List of Sample Output

show igmp interface on page 1789
show igmp interface brief on page 1790
show igmp interface detail on page 1790
show igmp interface <interface-name> on page 1790

Output Fields

Table 106 on page 1788 describes the output fields for the **show igmp interface** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2018, Juniper Networks, Inc.

Table 106: show igmp interface Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface.	All levels
Querier	Address of the routing device that has been elected to send membership queries.	All levels
State	State of the interface: **Up** or **Down**.	All levels
SSM Map Policy	Name of the source-specific multicast (SSM) map policy that has been applied to the IGMP interface.	All levels
Timeout	How long until the IGMP querier is declared to be unreachable, in seconds.	All levels
Version	IGMP version being used on the interface: **1** , **2** , or **3**.	All levels
Groups	Number of groups on the interface.	All levels
Group limit	Maximum number of groups allowed on the interface. Any joins requested after the limit is reached are rejected.	All levels
Group threshold	Configured threshold at which a warning message is generated. This threshold is based on a percentage of groups received on the interface. If the number of groups received reaches the configured threshold, the device generates a warning message.	All levels
Group log-interval	Time (in seconds) between consecutive log messages.	All levels
Immediate Leave	State of the immediate leave option: • **On**—Indicates that the router removes a host from the multicast group as soon as the router receives a leave group message from a host associated with the interface. • **Off**—Indicates that after receiving a leave group message, instead of removing a host from the multicast group immediately, the router sends a group query to determine if another receiver responds.	All levels
Promiscuous Mode	State of the promiscuous mode option: • **On**—Indicates that the router can accept IGMP reports from subnetworks that are not associated with its interfaces. • **Off**—Indicates that the router can accept IGMP reports only from subnetworks that are associated with its interfaces.	All levels
Distributed	State of IGMP, which, by default, takes place on the Routing Engine for MX Series routers but can be distributed to the Packet Forwarding Engine to provide faster processing of join and leave events. • **On**—distributed IGMP is enabled.	All levels

 Copyright © 2018, Juniper Networks, Inc.

Table 106: show igmp interface Output Fields (continued)

Field Name	Field Description	Level of Output
Passive	State of the passive mode option: • **On**—Indicates that the router can run IGMP on the interface but not send or receive control traffic such as IGMP reports, queries, and leaves. • **Off**—Indicates that the router can run IGMP on the interface and send or receive control traffic such as IGMP reports, queries, and leaves. The **passive** statement enables you to selectively activate up to two out of a possible three available query or control traffic options. When enabled, the following options appear after the **on** state declaration: • **send-general-query**—The interface sends general queries. • **send-group-query**—The interface sends group-specific and group-source-specific queries. • **allow-receive**—The interface receives control traffic.	All levels
OIF map	Name of the OIF map (if configured) associated with the interface.	All levels
SSM map	Name of the source-specific multicast (SSM) map (if configured) used on the interface.	All levels
Configured Parameters	Information configured by the user: • **IGMP Query Interval**—Interval (in seconds) at which this router sends membership queries when it is the querier. • **IGMP Query Response Interval**—Time (in seconds) that the router waits for a report in response to a general query. • **IGMP Last Member Query Interval**—Time (in seconds) that the router waits for a report in response to a group-specific query. • **IGMP Robustness Count**—Number of times the router retries a query.	All levels
Derived Parameters	Derived information: • **IGMP Membership Timeout**—Timeout period (in seconds) for group membership. If no report is received for these groups before the timeout expires, the group membership is removed. • **IGMP Other Querier Present Timeout**—Time (in seconds) that the router waits for the IGMP querier to send a query.	All levels

Sample Output

show igmp interface

```
user@host> show igmp interface
Interface: at-0/3/1.0
    Querier: 203.0.3.113.31
    State:         Up Timeout:    None Version:  2 Groups:      4
    SSM Map Policy: ssm-policy-A
Interface: so-1/0/0.0
    Querier: 203.0.113.11
    State:         Up Timeout:    None Version:  2 Groups:      2
    SSM Map Policy: ssm-policy-B
Interface: so-1/0/1.0
```

Copyright © 2018, Juniper Networks, Inc.

```
    Querier: 203.0.113.21
    State:         Up Timeout:    None Version:  2 Groups:      4
    SSM Map Policy: ssm-policy-C
Immediate Leave: On
Promiscuous Mode: Off
Passive: Off
Distributed: OnConfigured Parameters:

IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Derived Parameters:
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0
```

show igmp interface brief

The output for the **show igmp interface brief** command is identical to that for the **show igmp interface** command. For sample output, see show igmp interface on page 1789.

show igmp interface detail

The output for the **show igmp interface detail** command is identical to that for the **show igmp interface** command. For sample output, see show igmp interface on page 1789.

show igmp interface <interface-name>

```
user@host# show igmp interface ge-3/2/0.0
Interface: ge-3/2/0.0
  Querier: 203.0.113.111
  State: Up Timeout:    None
  Version:  3
  Groups:      1
  Group limit: 8
  Group threshold: 60
  Group log-interval: 10
  Immediate leave:  Off
  Promiscuous mode:  Off
  Distributed: On
```

 Copyright © 2018, Juniper Networks, Inc.

show igmp snooping interface

Syntax

show igmp snooping interface *interface-name*
<brief | detail>
<bridge-domain *bridge-domain-name*>
<logical-system *logical-system-name*>
<virtual-switch *virtual-switch-name*>
<vlan-id *vlan-identifier*>

Release Information

Command introduced in Junos OS Release 8.5.

Description

Display IGMP snooping interface information.

Options

none —Display detailed information.

brief | detail—(Optional) When applicable, this option lets you choose the how much detail to display.

bridge-domain *bridge-domain-name*—(Optional) Display information about a particular bridge domain.

logical-system *logical-system-name*—(Optional) Display information about a particular logical system, or type ’all’.

virtual-switch *virtual-switch-name*—(Optional) Display information about a particular virtual switch.

vlan-id *vlan-identifier*—(Optional) Display information about a particular VLAN.

Required Privilege Level

view

Related Documentation

- show igmp snooping membership on page 1796
- show igmp snooping statistics on page 1801

List of Sample Output

show igmp snooping interface on page 1792
show igmp snooping interface (logical systems) on page 1793
show igmp snooping interface (Group Limit Configured) on page 1795

Output Fields

Table 107 on page 1791 lists the output fields for the **show igmp snooping interface** command. Output fields are listed in the approximate order in which they appear.

Table 107: show igmp snooping interface Output Fields

Field Name	Field Description	Level of Output
Routing-instance	Routing instance for IGMP snooping.	All levels

Table 107: show igmp snooping interface Output Fields (continued)

Field Name	Field Description	Level of Output
Learning Domain	Learning domain for snooping.	All levels
IGMP Query Interval	Frequency (in seconds) with which this router sends membership queries when it is the querier.	All levels
IGMP Query Response Interval	Time (in seconds) that the router waits for a response to a general query.	All levels
IGMP Last Member Query Interval	Time (in seconds) that the router waits for a report in response to a group-specific query.	All levels
IGMP Robustness Count	Number of times the router retries a query.	All levels
immediate-leave	State of immediate leave: **On** or **Off**.	All levels
router-interface	Router interfaces that are part of this learning domain.	All levels
Group limit	Maximum number of (source,group) pairs allowed per interface. When a group limit is not configured, this field is not shown.	All levels
interface	Interfaces that are being snooped in this learning domain.	All levels
Groups	Number of groups on the interface.	All levels
State	State of the interface: **Up** or **Down**.	All levels
Up Groups	Number of active multicast groups attached to the logical interface.	All levels
IGMP Membeship Timeout	Timeout for group membership. If no report is received for these groups before the timeout expires, the group membership is removed.	All levels
IGMP Other Querier Present Timeout	Time that the router waits for the IGMP querier to send a query.	All levels

Sample Output

show igmp snooping interface

```
user@host> show igmp snooping interface ge-0/1/4
Instance: default-switch

Bridge-Domain: sample

Learning-Domain: default
Interface: ge-0/1/4.0
State: Up Groups: 0
Immediate leave: Off
Router interface: no
```

 Copyright © 2018, Juniper Networks, Inc.

```
Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Derived Parameters:
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0
```

show igmp snooping interface (logical systems)

```
user@host> show igmp snooping interface logical-system all
logical-system: default
Instance: VPLS-6
Learning-Domain: default
Interface: ge-0/2/2.601
    State:         Up Groups:     10
    Immediate leave: Off
    Router interface: no

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Instance: VS-4
Bridge-Domain: VS-4-BD-1
Learning-Domain: vlan-id 1041
Interface: ae2.3
    State:         Up Groups:      0
    Immediate leave: Off
    Router interface: no
Interface: ge-0/2/2.1041
    State:         Up Groups:     20
    Immediate leave: Off
    Router interface: no

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Instance: default-switch
Bridge-Domain: bd-200
Learning-Domain: default
Interface: ge-0/2/2.100
    State:         Up Groups:     20
    Immediate leave: Off
    Router interface: no

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Bridge-Domain: bd0
```

Copyright © 2018, Juniper Networks, Inc.

```
Learning-Domain: default
Interface: ae0.0
    State:         Up Groups:      0
    Immediate leave: Off
    Router interface: yes
Interface: ae1.0
    State:         Up Groups:      0
    Immediate leave: Off
    Router interface: no
Interface: ge-0/2/2.0
    State:         Up Groups:     32
    Immediate leave: Off
    Router interface: no

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Instance: VPLS-1
Learning-Domain: default
Interface: ge-0/2/2.502
    State:         Up Groups:     11
    Immediate leave: Off
    Router interface: no

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Instance: VS-1
Bridge-Domain: VS-BD-1
Learning-Domain: default
Interface: ae2.0
    State:         Up Groups:      0
    Immediate leave: Off
    Router interface: no
Interface: ge-0/2/2.1010
    State:         Up Groups:     20
    Immediate leave: Off
    Router interface: no

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Bridge-Domain: VS-BD-2
Learning-Domain: default
Interface: ae2.0
    State:         Up Groups:      0
    Immediate leave: Off
    Router interface: no
Interface: ge-0/2/2.1011
    State:         Up Groups:     20
    Immediate leave: Off
    Router interface: no
```

 Copyright © 2018, Juniper Networks, Inc.

```
Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Instance: VPLS-p2mp
Learning-Domain: default
Interface: ge-0/2/2.3001
    State:         Up Groups:      0
    Immediate leave: Off
    Router interface: no

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2
```

show igmp snooping interface (Group Limit Configured)

```
user@host> show igmp snooping interface instance vpls1
Instance: vpls1

Learning-Domain: default
Interface: ge-1/3/9.0
    State:         Up Groups:      0
    Immediate leave: Off
    Router interface: yes
Interface: ge-1/3/8.0
    State:         Up Groups:      0
    Immediate leave: Off
    Router interface: yes
    Group limit:   1000

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2
```

Copyright © 2018, Juniper Networks, Inc.

show igmp snooping membership

Syntax

show igmp snooping membership
<brief | detail>
<interface *interface-name*>
<vlan (*vlan-id* | *vlan-name*)>
<bridge-domain *bridge-domain-name*>
<group *group-name*>
<logical-system *logical-system-name*>
<virtual-switch *virtual-switch-name*>
<vlan-id *vlan-identifier*>

Release Information

Command introduced in Junos OS Release 8.5.
Command introduced in Junos OS Release 18.1R1 for the SRX1500 devices.

Description

Display the multicast group membership information maintained by IGMP snooping.

Options

none—Display the multicast group membership information about all VLANs on which IGMP snooping is enabled.

brief | detail—(Optional) Display the specified level of output. The default is **brief**.

interface *interface-name*—(Optional) Display the multicast group membership information about the specified interface.

vlan (*vlan-id* | *vlan-name*)—(Optional) Display the multicast group membership for the specified VLAN.

bridge-domain *bridge-domain-name*—(Optional) Display information about a particular bridge domain.

group *group-name* —(Optional) Display information about this group address.

logical-system *logical-system-name*—(Optional) Display information about a particular logical system, or type 'all'.

virtual-switch *virtual-switch-name*—(Optional) Display information about a particular virtual switch.

vlan-id *vlan-identifier*—(Optional) Display information about a particular VLAN.

Required Privilege Level

view

Related Documentation

- show igmp snooping interface on page 1791
- show igmp snooping statistics on page 1801
- clear igmp snooping membership on page 1732

Copyright © 2018, Juniper Networks, Inc.

List of Sample Output

Output Fields

Table 108 on page 1797 lists the output fields for the **show igmp snooping membership** command. Output fields are listed in the approximate order in which they appear.

Table 108: show igmp snooping membership Output Fields

Field Name	Field Description	Level of Output
VLAN	Name of the VLAN.	All
Instance	Routing instance for IGMP snooping.	All levels
Learning Domain	Learning domain for snooping.	All levels
Interface	Interface on which this router is a proxy.	**detail**
Up Groups	Number of active multicast groups attached to the logical interface.	All levels
Group	Multicast group address in the membership database.	All levels
Group Mode	Mode the SSM group is operating in: **Include** or **Exclude**.	All levels
Source	Source address used on queries.	**detail**
Last reported by	Address of source last replying to the query.	**detail**
Group Timeout	Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer.	All levels
Timeout	Length of time (in seconds) left until the entry is purged.	**detail**
Type	Way that the group membership information was learned: • **Dynamic**—Group membership was learned by the IGMP protocol. • **Static**—Group membership was learned by configuration.	**detail**
Include receiver	Source address of receiver included in membership with timeout (in seconds).	**detail**

Sample Output

show igmp snooping membership

```
user@host> show igmp snooping membership
Instance: vpls2

Learning-Domain: vlan-id 2
Interface: ge-3/0/0.2
Up Groups:      0
Interface: ge-3/1/0.2
Up Groups:      0
Interface: ge-3/1/5.2
Up Groups:      0

Instance: vpls1

Learning-Domain: vlan-id 1
Interface: ge-3/0/0.1
Up Groups:      0
Interface: ge-3/1/0.1
Up Groups:      0
Interface: ge-3/1/5.1
Up Groups:      1
    Group: 233.252.0.99
        Group mode: Exclude
        Source: 0.0.0.0
        Last reported by: 233.252.0.87
        Group timeout:     173 Type: Dynamic
```

show igmp-snooping membership (SRX1500)

```
user@host> show igmp-snooping membership
Instance: default-switch

Vlan: v1

Learning-Domain: default
Interface: ge-0/0/3.0, Groups: 1
Group: 233.252.0.100
Group mode: Exclude
Source: 0.0.0.0
Last reported by: Local
Group timeout: 0 Type: Static
```

show igmp-snooping membership detail (SRX1500)

```
user@host> show igmp-snooping membership detail

VLAN: vlan2 Tag: 2 (Index: 3)
    Router interfaces:
        ge-1/0/0.0 dynamic Uptime: 00:14:24 timeout: 253
    Group: 233.252.0.99
      ge-1/0/17.0 259 Last reporter: 10.0.0.90 Receiver count: 1
      Uptime: 00:00:19 timeout: 259 Flags: <V3-hosts>
        Include source: 10.2.11.5, 10.2.11.12
```

 Copyright © 2018, Juniper Networks, Inc.

show igmp snooping membership (Exclude Mode)

```
user@host> show igmp snooping membership
Instance: vpls2

Learning-Domain: vlan-id 2
Interface: ge-3/0/0.2
Up Groups:      0
Interface: ge-3/1/0.2
Up Groups:      0
Interface: ge-3/1/5.2
Up Groups:      0

Instance: vpls1

Learning-Domain: vlan-id 1
Interface: ge-3/0/0.1
Up Groups:      0
Interface: ge-3/1/0.1
Up Groups:      0
Interface: ge-3/1/5.1
Up Groups:      1
    Group: 233.252.0.99
        Group mode: Exclude
        Source: 0.0.0.0
        Last reported by: 233.252.0.87
        Group timeout:     173 Type: Dynamic
```

show igmp-snooping membership detail (SRX1500)

```
user@host> show igmp-snooping membership detail

VLAN: vlan2 Tag: 2 (Index: 3)
    Router interfaces:
        ge-1/0/0.0 dynamic Uptime: 00:14:24 timeout: 253
    Group: 233.252.0.99
      ge-1/0/17.0 259 Last reporter: 233.252.0.82 Receiver count: 1
      Uptime: 00:00:19 timeout: 259 Flags: <V3-hosts>
        Include source: 233.252.0.84, 233.252.0.83
```

show igmp-snooping membership vlan detail (SRX1500)

```
user@host> show igmp-snooping membership vlan vlan700 detail
VLAN: vlan700 Tag: 700 (Index: 52)
    Router interfaces:
        ae2.0 dynamic Uptime: 16:53:13 timeout: 245
   Group: 233.252.0.1
50      ge-0/0/1.0 Last reporter: 233.252.0.87
     Uptime: 17:00:52 timeout: 237 Flags: <V2-hosts>
     ge-0/0/0.0 Last reporter: 10.2.188.202
     Uptime: 17:00:50 timeout: 243 Flags: <V2-hosts>
```

show igmp snooping membership interface ge-0/1/2.200

```
user@host> show igmp snooping membership interface ge-0/1/2.200
Instance: bridge-domain bar

Learning-Domain: default
Interface: ge-0/1/2.200
```

Copyright © 2018, Juniper Networks, Inc.

```
    Group: 233.252.0.1
        Source: 0.0.0.0
        Timeout:     391 Type: Static
    Group: 232.1.1.1
        Source: 192.128.1.1
        Timeout:       0 Type: Static
```

show igmp snooping membership vlan-id 1

```
user@host> show igmp snooping membership vlan–id 1
Instance: vpls2

Instance: vpls1

Learning-Domain: vlan-id 1
Interface: ge-3/0/0.1
Up Groups:      0
Interface: ge-3/1/0.1
Up Groups:      0
Interface: ge-3/1/5.1
Up Groups:      1
    Group: 233.252.0.1
        Group mode: Exclude
        Source: 0.0.0.0
        Last reported by: 233.252.0.82
        Group timeout:     209 Type: Dynamic
```

 Copyright © 2018, Juniper Networks, Inc.

show igmp snooping statistics

Syntax

show igmp snooping statistics
<brief | detail>
<bridge-domain *bridge-domain-name*>
<logical-system *logical-system-name*>
<virtual-switch *virtual-switch-name*>
<vlan-id *vlan-identifier*>

Release Information

Command introduced in Junos OS Release 8.5.
Command introduced in Junos OS Release 18.1R1 for the SR1500 devices.

Description

Display IGMP snooping statistics.

Options

none—(Optional) Display detailed information.

brief | detail—(Optional) Display the specified level of output.

bridge-domain ***bridge-domain-name***—(Optional) Display information about a particular bridge domain.

logical-system ***logical-system-name***—(Optional) Display information about a particular logical system, or type 'all'.

virtual-switch ***virtual-switch-name***—(Optional) Display information about a particular virtual switch.

vlan-id ***vlan-identifier***—(Optional) Display information about a particular VLAN.

Required Privilege Level

view

Related Documentation

- show igmp snooping interface on page 1791
- show igmp snooping membership on page 1796
- clear igmp snooping statistics on page 1734

List of Sample Output

show igmp snooping statistics on page 1803
show igmp-snooping statistics (SRX1500) on page 1803
show igmp snooping statistics logical-systems all on page 1804
show igmp snooping statistics interface (Bridge Domains Configured) on page 1805

Output Fields

Table 109 on page 1802 lists the output fields for the **show igmp snooping statistics** command. Output fields are listed in the approximate order in which they appear.

Table 109: show igmp snooping statistics Output Fields

Field Name	Field Description	Level of Output
Routing-instance	Routing instance for IGMP snooping.	All levels
IGMP packet statistics	Heading for IGMP snooping statistics for all interfaces or for the specified interface.	All levels
learning-domain	Appears at end of “IGMP packets statistics” line.	All levels
IGMP Message type	Summary of IGMP statistics: • **Membership Query**—Number of membership queries sent and received. • **V1 Membership Report**—Number of version 1 membership reports sent and received. • **DVMRP**—Number of DVMRP messages sent or received. • **PIM V1**—Number of PIM version 1 messages sent or received. • **Cisco Trace**—Number of Cisco trace messages sent or received. • **V2 Membership Report**—Number of version 2 membership reports sent or received. • **Group Leave**—Number of group leave messages sent or received. • **Domain Wide Report**—Number of domain-wide reports sent or received. • **V3 Membership Report**—Number of version 3 membership reports sent or received. • **Other Unknown types**—Number of unknown message types received. • **IGMP v3 unsupported type**—Number of messages received with unknown and unsupported IGMP version 3 message types. • **IGMP v3 source required for SSM**—Number of IGMP version 3 messages received that contained no source. • **IGMP v3 mode not applicable for SSM**—Number of IGMP version 3 messages received that did not contain a mode applicable for source-specific multicast (SSM).	All levels
Received	Number of messages received.	All levels
Sent	Number of messages sent.	All levels
Rx errors	Number of received packets that contained errors.	All levels
IGMP Global Statistics	Summary of IGMP snooping statistics for all interfaces. • **Bad Length**—Number of messages received with length errors so severe that further classification could not occur. • **Bad Checksum**—Number of messages received with a bad IP checksum. No further classification was performed. • **Rx non-local**—Number of messages received from senders that are not local.	All levels

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show igmp snooping statistics

```
user@host> show igmp snooping statistics
Routing-instance foo

 IGMP packet statistics for all interfaces in learning-domain vlan-100

 IGMP Message type      Received       Sent  Rx errors
 Membership Query             89         51          0
 V1 Membership Report          0          0          0
 DVMRP                         0          0          0
 PIM V1                        0          0          0
 Cisco Trace                   0          0          0
 V2 Membership Report        139          0          0
 Group Leave                   0          0          0
 Domain Wide Report            0          0          0
 V3 Membership Report        136          0          0
 Other Unknown types                                 0
 IGMP v3 unsupported type                            0
 IGMP v3 source required for SSM                    23
 IGMP v3 mode not applicable for SSM                 0

 IGMP Global Statistics
 Bad Length                    0
 Bad Checksum                  0
 Rx non-local                  0

Routing-instance bar

 IGMP packet statistics for all interfaces in learning-domain vlan-100

 IGMP Message type      Received       Sent  Rx errors
 Membership Query             89         51          0
 V1 Membership Report          0          0          0
 DVMRP                         0          0          0
 PIM V1                        0          0          0
 Cisco Trace                   0          0          0
 V2 Membership Report        139          0          0
 Group Leave                   0          0          0
 Domain Wide Report            0          0          0
 V3 Membership Report        136          0          0
 Other Unknown types                                 0
 IGMP v3 unsupported type                            0
 IGMP v3 source required for SSM                    23
 IGMP v3 mode not applicable for SSM                 0

 IGMP Global Statistics
 Bad Length                    0
 Bad Checksum                  0
 Rx non-local                  0
```

show igmp-snooping statistics (SRX1500)

```
user@host> show igmp-snooping statistics
Vlan: v1
IGMP Message type      Received       Sent  Rx errors
Membership Query              0          0          0
V1 Membership Report          0          0          0
```

Copyright © 2018, Juniper Networks, Inc.

```
DVMRP                        0          0          0
PIM V1                       0          0          0
Cisco Trace                  0          0          0
V2 Membership Report         0          0          0
Group Leave                  0          0          0
Mtrace Response              0          0          0
Mtrace Request               0          0          0
Domain Wide Report           0          0          0
V3 Membership Report         0          0          0
Other Unknown types                                0
```

show igmp snooping statistics logical-systems all

```
user@host> show igmp snooping statistics logical-systems all

logical-system: default
Bridge: VPLS-6
IGMP Message type       Received       Sent  Rx errors
Membership Query               0          4          0
V1 Membership Report           0          0          0
DVMRP                          0          0          0
PIM V1                         0          0          0
Cisco Trace                    0          0          0
V2 Membership Report           0          0          0
Group Leave                    0          0          0
Mtrace Response                0          0          0
Mtrace Request                 0          0          0
Domain Wide Report             0          0          0
V3 Membership Report           0          0          0
Other Unknown types                                  0

Learning-Domain: vlan-id 1041 bridge-domain VS-4-BD-1
IGMP Message type       Received       Sent  Rx errors
Membership Query               0          4          0
V1 Membership Report           0          0          0
DVMRP                          0          0          0
PIM V1                         0          0          0
Cisco Trace                    0          0          0
V2 Membership Report           0          0          0
Group Leave                    0          0          0
Mtrace Response                0          0          0
Mtrace Request                 0          0          0
Domain Wide Report             0          0          0
V3 Membership Report           0          0          0
Other Unknown types                                  0

Bridge: VPLS-p2mp
IGMP Message type       Received       Sent  Rx errors
Membership Query               0          2          0
V1 Membership Report           0          0          0
DVMRP                          0          0          0
PIM V1                         0          0          0
Cisco Trace                    0          0          0
V2 Membership Report           0          0          0
Group Leave                    0          0          0
Mtrace Response                0          0          0
Mtrace Request                 0          0          0
Domain Wide Report             0          0          0
V3 Membership Report           0          0          0
Other Unknown types                                  0
```


Copyright © 2018, Juniper Networks, Inc.

```
Bridge: VS-BD-1
IGMP Message type      Received       Sent  Rx errors
Membership Query              0          6          0
V1 Membership Report          0          0          0
DVMRP                         0          0          0
PIM V1                        0          0          0
Cisco Trace                   0          0          0
V2 Membership Report          0          0          0
Group Leave                   0          0          0
Mtrace Response               0          0          0
Mtrace Request                0          0          0
Domain Wide Report            0          0          0
V3 Membership Report          0          0          0
Other Unknown types                                 0
```

show igmp snooping statistics interface (Bridge Domains Configured)

user@host> **show igmp snooping statistics interface**

```
Bridge: bridge-domain1
IGMP interface packet statistics for ge-2/0/8.0
IGMP Message type      Received       Sent  Rx errors
Membership Query              0          2          0
V1 Membership Report          0          0          0
DVMRP                         0          0          0
PIM V1                        0          0          0
Cisco Trace                   0          0          0
V2 Membership Report          0          0          0
Group Leave                   0          0          0
Mtrace Response               0          0          0
Mtrace Request                0          0          0
Domain Wide Report            0          0          0
V3 Membership Report          0          0          0
Other Unknown types                                 0

Bridge: bridge-domain2
IGMP interface packet statistics for ge-2/0/8.0
IGMP Message type      Received       Sent  Rx errors
Membership Query              0          2          0
V1 Membership Report          0          0          0
DVMRP                         0          0          0
PIM V1                        0          0          0
Cisco Trace                   0          0          0
V2 Membership Report          0          0          0
Group Leave                   0          0          0
Mtrace Response               0          0          0
Mtrace Request                0          0          0
Domain Wide Report            0          0          0
V3 Membership Report          0          0          0
Other Unknown types                                 0
```

Copyright © 2018, Juniper Networks, Inc.

show igmp statistics

List of Syntax

Syntax on page 1806
Syntax (EX Series Switch and the QFX Series) on page 1806

Syntax

show igmp statistics
<brief | detail>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)

show igmp statistics
<brief | detail>
<interface *interface-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display Internet Group Management Protocol (IGMP) statistics.

Options

none—Display IGMP statistics for all interfaces.

brief | detail—(Optional) Display the specified level of output.

interface *interface-name*—(Optional) Display IGMP statistics about the specified interface only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- clear igmp statistics on page 1735

List of Sample Output

show igmp statistics on page 1808
show igmp statistics interface on page 1808

Output Fields

Table 110 on page 1806 describes the output fields for the **show igmp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 110: show igmp statistics Output Fields

Field Name	Field Description
IGMP packet statistics	Heading for IGMP packet statistics for all interfaces or for the specified interface name.

Copyright © 2018, Juniper Networks, Inc.

Table 110: show igmp statistics Output Fields (continued)

Field Name	Field Description
IGMP Message type	Summary of IGMP statistics: • **Membership Query**—Number of membership queries sent and received. • **V1 Membership Report**—Number of version 1 membership reports sent and received. • **DVMRP**—Number of DVMRP messages sent or received. • **PIM V1**—Number of PIM version 1 messages sent or received. • **Cisco Trace**—Number of Cisco trace messages sent or received. • **V2 Membership Report**—Number of version 2 membership reports sent or received. • **Group Leave**—Number of group leave messages sent or received. • **Mtrace Response**—Number of Mtrace response messages sent or received. • **Mtrace Request**—Number of Mtrace request messages sent or received. • **Domain Wide Report**—Number of domain-wide reports sent or received. • **V3 Membership Report**—Number of version 3 membership reports sent or received. • **Other Unknown types**—Number of unknown message types received. • **IGMP v3 unsupported type**—Number of messages received with unknown and unsupported IGMP version 3 message types. • **IGMP v3 source required for SSM**—Number of IGMP version 3 messages received that contained no source. • **IGMP v3 mode not applicable for SSM**—Number of IGMP version 3 messages received that did not contain a mode applicable for source-specific multicast (SSM). Beginning with certain releases, this type includes records received for groups in the SSM range of addresses and in which the mode is MODE_IS_EXCLUDE or CHANGE_TO_EXCLUDE_MODE. This includes records with a non-empty source list.
Received	Number of messages received.
Sent	Number of messages sent.
Rx errors	Number of received packets that contained errors.
Max Rx rate (pps)	Maximum number of IGMP packets received during 1 second interval.
IGMP Global Statistics	Summary of IGMP statistics for all interfaces. • **Bad Length**—Number of messages received with length errors so severe that further classification could not occur. • **Bad Checksum**—Number of messages received with a bad IP checksum. No further classification was performed. • **Bad Receive If**—Number of messages received on an interface not enabled for IGMP. • **Rx non-local**—Number of messages received from senders that are not local. • **Timed out**—Number of groups that timed out as a result of not receiving an explicit leave message. • **Rejected Report**—Number of reports dropped because of the IGMP group policy. • **Total Interfaces**—Number of interfaces configured to support IGMP.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show igmp statistics

```
user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type       Received        Sent  Rx errors
Membership Query            8883         459          0
V1 Membership Report           0           0          0
DVMRP                          0           0          0
PIM V1                         0           0          0
Cisco Trace                    0           0          0
V2 Membership Report           0           0          0
Group Leave                    0           0          0
Mtrace Response                0           0          0
Mtrace Request                 0           0          0
Domain Wide Report             0           0          0
V3 Membership Report           0           0          0
Other Unknown types                                   0
IGMP v3 unsupported type                              0
IGMP v3 source required for SSM                       0
IGMP v3 mode not applicable for SSM                   0

IGMP Global Statistics
Bad Length                     0
Bad Checksum                   0
Bad Receive If                 0
Rx non-local                1227
Timed out                      0
Rejected Report                0
Total Interfaces               2
Max Rx rate (pps)           1536
```

show igmp statistics interface

```
user@host> show igmp statistics interface fe-1/0/1.0
IGMP interface packet statistics for fe-1/0/1.0
IGMP Message type       Received        Sent  Rx errors
Membership Query               0         230          0
V1 Membership Report           0           0          0
```


Copyright © 2018, Juniper Networks, Inc.

show mld group

Syntax

show mld group
<brief | detail>
<*group-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display information about Multicast Listener Discovery (MLD) group membership.

Options

none—Display standard information about all MLD groups.

brief | detail—(Optional) Display the specified level of output.

group-name—(Optional) Display MLD information about the specified group.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

• clear mld membership on page 1737

List of Sample Output

show mld group (Include Mode) on page 1810
show mld group (Exclude Mode) on page 1811
show mld group brief on page 1811
show mld group detail (Include Mode) on page 1811
show mld group detail (Exclude Mode) on page 1812

Output Fields

Table 111 on page 1809 describes the output fields for the **show mld group** command. Output fields are listed in the approximate order in which they appear.

Table 111: show mld group Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface that received the MLD membership report; **local** means that the local router joined the group itself.	All levels
Group	Group address.	All levels
Source	Source address.	All levels
Group Mode	Mode the SSM group is operating in: **Include** or **Exclude**.	All levels
Last reported by	Address of the host that last reported membership in this group.	All levels

Copyright © 2018, Juniper Networks, Inc.

Table 111: show mld group Output Fields (continued)

Field Name	Field Description	Level of Output
Source timeout	Time remaining until the group traffic is no longer forwarded. The timer is refreshed when a listener in include mode sends a report. A group in exclude mode or configured as a static group displays a zero timer.	**detail**
Timeout	Time remaining until the group membership is removed.	**brief none**
Group timeout	Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer.	**detail**
Type	Type of group membership: • **Dynamic**—Host reported the membership. • **Static**—Membership is configured.	All levels

Sample Output

show mld group (Include Mode)

```
user@host> show mld group
Interface: fe-0/1/2.0
    Group: ff02::1:ff05:1a67
        Group mode: Include
        Source: ::
        Last reported by: fe80::2e0:81ff:fe05:1a67
        Timeout:     245 Type: Dynamic
    Group: ff02::1:ffa8:c35e
        Group mode: Include
        Source: ::
        Last reported by: fe80::2e0:81ff:fe05:1a67
        Timeout:     241 Type: Dynamic
    Group: ff02::2:43e:d7f6
        Group mode: Include
        Source: ::
        Last reported by: fe80::2e0:81ff:fe05:1a67
        Timeout:     244 Type: Dynamic
    Group: ff05::2
        Group mode: Include
        Source: ::
        Last reported by: fe80::2e0:81ff:fe05:1a67
        Timeout:     244 Type: Dynamic
Interface: local
    Group: ff02::2
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: ff02::16
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
```

 Copyright © 2018, Juniper Networks, Inc.

show mld group (Exclude Mode)

```
user@host> show mld group
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
    Group: ff02::6
        Source: ::
        Last reported by: fe80::21f:12ff:feb6:4b3a
        Timeout:     245 Type: Dynamic
    Group: ff02::16
        Source: ::
        Last reported by: fe80::21f:12ff:feb6:4b3a
        Timeout:      28 Type: Dynamic
Interface: local
    Group: ff02::2
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: ff02::16
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
```

show mld group brief

The output for the **show mld group brief** command is identical to that for the **show mld group** command. For sample output, see show mld group (Include Mode) on page 1810 show mld group (Exclude Mode) on page 1811.

show mld group detail (Include Mode)

```
user@host> show mld group detail
Interface: fe-0/1/2.0
    Group: ff02::1:ff05:1a67
        Group mode: Include
        Source: ::
        Last reported by: fe80::2e0:81ff:fe05:1a67
        Timeout:     224 Type: Dynamic
    Group: ff02::1:ffa8:c35e
        Group mode: Include
        Source: ::
        Last reported by: fe80::2e0:81ff:fe05:1a67
        Timeout:     220 Type: Dynamic
    Group: ff02::2:43e:d7f6
        Group mode: Include
        Source: ::
        Last reported by: fe80::2e0:81ff:fe05:1a67
        Timeout:     223 Type: Dynamic
    Group: ff05::2
        Group mode: Include
        Source: ::
        Last reported by: fe80::2e0:81ff:fe05:1a67
        Timeout:     223 Type: Dynamic
Interface: so-1/0/1.0
    Group: ff02::2
        Group mode: Include
        Source: ::
        Last reported by: fe80::280:42ff:fe15:f445
        Timeout:     258 Type: Dynamic
```

Copyright © 2018, Juniper Networks, Inc.

```
Interface: local
    Group: ff02::2
        Group mode: Include
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: ff02::16
        Source: ::
        Last reported by: Local
        Timeout:       0 Type: Dynamic
```

show mld group detail (Exclude Mode)

```
user@host> show mld group detail
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
    Group: ff02::6
        Group mode: Exclude
        Source: ::
        Source timeout: 0
        Last reported by: fe80::21f:12ff:feb6:4b3a
        Group timeout:     226 Type: Dynamic
    Group: ff02::16
        Group mode: Exclude
        Source: ::
        Source timeout: 0
        Last reported by: fe80::21f:12ff:feb6:4b3a
        Group timeout:     246 Type: Dynamic
Interface: local
    Group: ff02::2
        Group mode: Exclude
        Source: ::
        Source timeout: 0
        Last reported by: Local
        Group timeout:       0 Type: Dynamic
    Group: ff02::16
        Group mode: Exclude
        Source: ::
        Source timeout: 0
        Last reported by: Local
        Group timeout:       0 Type: Dynamic
```


Copyright © 2018, Juniper Networks, Inc.

show mld interface

Syntax

show mld interface
<brief | detail>
<*interface-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display information about multipoint Listener Discovery (MLD)-enabled interfaces.

Options

none—Display standard information about all MLD-enabled interfaces.

brief | detail—(Optional) Display the specified level of output.

interface-name—(Optional) Display information about the specified interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

• clear mld membership on page 1737

List of Sample Output

show mld interface on page 1815
show mld interface brief on page 1816
show mld interface detail on page 1816
show mld interface <interface-name> on page 1816

Output Fields

Table 112 on page 1813 describes the output fields for the **show mld interface** command. Output fields are listed in the approximate order in which they appear.

Table 112: show mld interface Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface.	All levels
Querier	Address of the router that has been elected to send membership queries.	All levels
State	State of the interface: **Up** or **Down**.	All levels
SSM Map Policy	Name of the source-specific multicast (SSM) map policy that has been applied to the interface.	All levels
SSM Map Policy	Name of the source-specific multicast (SSM) map policy at the MLD interface.	All levels

Copyright © 2018, Juniper Networks, Inc.

Table 112: show mld interface Output Fields (continued)

Field Name	Field Description	Level of Output
Timeout	How long until the MLD querier is declared to be unreachable, in seconds.	All levels
Version	MLD version being used on the interface: **1** or **2**.	All levels
Groups	Number of groups on the interface.	All levels
Passive	State of the passive mode option: • **On**—Indicates that the router can run IGMP or MLD on the interface but not send or receive control traffic such as IGMP or MLD reports, queries, and leaves. • **Off**—Indicates that the router can run IGMP or MLD on the interface and send or receive control traffic such as IGMP or MLD reports, queries, and leaves. The **passive** statement enables you to selectively activate up to two out of a possible three available query or control traffic options. When enabled, the following options appear after the **on** state declaration: • **send-general-query**—The interface sends general queries. • **send-group-query**—The interface sends group-specific and group-source-specific queries. • **allow-receive**—The interface receives control traffic	All levels
OIF map	Name of the OIF map associated to the interface.	All levels
SSM map	Name of the source-specific multicast (SSM) map used on the interface, if configured.	All levels
Group limit	Maximum number of groups allowed on the interface. Any memberships requested after the limit is reached are rejected.	All levels
Group threshold	Configured threshold at which a warning message is generated. This threshold is based on a percentage of groups received on the interface. If the number of groups received reaches the configured threshold, the device generates a warning message.	All levels
Group log-interval	Time (in seconds) between consecutive log messages.	All levels
Immediate Leave	State of the immediate leave option: • **On**—Indicates that the router removes a host from the multicast group as soon as the router receives a multicast listener done message from a host associated with the interface. • **Off**—Indicates that after receiving a multicast listener done message, instead of removing a host from the multicast group immediately, the router sends a group query to determine if another receiver responds.	All levels

 Copyright © 2018, Juniper Networks, Inc.

Table 112: show mld interface Output Fields (continued)

Field Name	Field Description	Level of Output
Distributed	State of MLD, which, by default, takes place on the Routing Engine for MX Series routers but can be distributed to the Packet Forwarding Engine to provide faster processing of join and leave events. • **On**—distributed MLD is enabled.	All levels
Configured Parameters	Information configured by the user. • **MLD Query Interval (.1 secs)**—Interval at which this router sends membership queries when it is the querier. • **MLD Query Response Interval (.1 secs)**—Time that the router waits for a report in response to a general query. • **MLD Last Member Query Interval (.1 secs)**—Time that the router waits for a report in response to a group-specific query. • **MLD Robustness Count**—Number of times the router retries a query.	All levels
Derived Parameters	Derived information. • **MLD Membership Timeout (.1 secs)**—Timeout period for group membership. If no report is received for these groups before the timeout expires, the group membership will be removed. • **MLD Other Querier Present Timeout (.1 secs)**—Time that the router waits for the IGMP querier to send a query.	All levels

Sample Output

show mld interface

```
user@host> show mld interface
Interface: fe-0/0/0
    Querier: None
    State: Up         Timeout:       0    Version:  1    Groups:      0
    SSM Map Policy: ssm-policy-A
Interface: at-0/3/1.0
    Querier:  8038::c0a8:c345
    State: Up         Timeout:    None    Version:  1    Groups:      0
    SSM Map Policy: ssm-policy-B
Interface: fe-1/0/1.0
    Querier:  ::192.168.195.73
    State: Up         Timeout:    None    Version:  1    Groups:      3
    SSM Map Policy: ssm-policy-C
    SSM map: ipv6map1
Immediate Leave: On

Promiscuous Mode: Off
Passive: Off
Distributed: OnConfigured Parameters:

Configured Parameters:
MLD Query Interval (.1 secs): 1250
MLD Query Response Interval (.1 secs): 100
MLD Last Member Query Interval (.1 secs): 10
MLD Robustness Count: 2
```

Copyright © 2018, Juniper Networks, Inc.

```
Derived Parameters:
MLD Membership Timeout (.1secs): 2600
MLD Other Querier Present Timeout (.1 secs): 2550
```

show mld interface brief

The output for the **show mld interface brief** command is identical to that for the **show mld interface** command. For sample output, see show mld interface on page 1815.

show mld interface detail

The output for the **show mld interface detail** command is identical to that for the **show mld interface** command. For sample output, see show mld interface on page 1815.

show mld interface <interface-name>

```
user@host# show mld interface ge-3/2/0.0
Interface: ge-3/2/0.0
  Querier: 203.0.113.111
  State: Up Timeout:    None Version:  3 Groups:      1
  Group limit: 8
  Group threshold: 60
  Group log-interval: 10
  Immediate leave:  Off
  Promiscuous mode:  Off   Distributed: On
```


Copyright © 2018, Juniper Networks, Inc.

show mld statistics

Syntax

show mld statistics
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display information about Multicast Listener Discovery (MLD) statistics.

Options

none—Display MLD statistics for all interfaces.

interface *interface-name*—(Optional) Display statistics about the specified interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- clear mld statistics on page 1738

List of Sample Output

show mld statistics on page 1818
show mld statistics interface on page 1819

Output Fields

Table 113 on page 1817 describes the output fields for the **show mld statistics** command. Output fields are listed in the approximate order in which they appear.

Table 113: show mld statistics Output Fields

Field Name	Field Description
Received	Number of received packets.
Sent	Number of transmitted packets.
Rx errors	Number of received packets that contained errors.

Table 113: show mld statistics Output Fields (continued)

Field Name	Field Description
MLD Message type	Summary of MLD statistics. • **Listener Query (v1/v2)**—Number of membership queries sent and received. • **Listener Report (v1)**—Number of version 1 membership reports sent and received. • **Listener Done (v1/v2)**—Number of Listener Done messages sent and received. • **Listener Report (v2)**—Number of version 2 membership reports sent and received. • **Other Unknown types**—Number of unknown message types received. • **MLD v2 source required for SSM**—Number of MLD version 2 messages received that contained no source. • **MLD v2 mode not applicable for SSM**—Number of MLD version 2 messages received that did not contain a mode applicable for source-specific multicast (SSM).
MLD Global Statistics	Summary of MLD statistics for all interfaces. • **Bad Length**—Number of messages received with length errors so severe that further classification could not occur. • **Bad Checksum**—Number of messages received with an invalid IP checksum. No further classification was performed. • **Bad Receive If**—Number of messages received on an interface not enabled for MLD. • **Rx non-local**—Number of messages received from nonlocal senders. • **Timed out**—Number of groups that timed out as a result of not receiving an explicit leave message. • **Rejected Report**—Number of reports dropped because of the MLD group policy. • **Total Interfaces**—Number of interfaces configured to support IGMP.

Sample Output

show mld statistics

```
user@host> show mld statistics
MLD packet statistics for all interfaces
MLD Message type            Received       Sent  Rx errors
Listener Query (v1/v2)             0          2          0
Listener Report (v1)               0          0          0
Listener Done (v1/v2)              0          0          0
Listener Report (v2)               0          0          0
Other Unknown types                                      0
MLD v2 source required for SSM     2
MLD v2 mode not applicable for SSM 0

MLD Global Statistics
Bad Length                         0
Bad Checksum                       0
Bad Receive If                     0
Rx non-local                       0
```

 Copyright © 2018, Juniper Networks, Inc.

```
Timed out                          0
Rejected Report                    0
Total Interfaces                   2
```

show mld statistics interface

```
user@host> show mld statistics interface fe-1/0/1.0
MLD interface packet statistics for fe-1/0/1.0
MLD Message type            Received        Sent  Rx errors
Listener Query (v1/v2)             0           2          0
Listener Report (v1)               0           0          0
Listener Done (v1/v2)              0           0          0
Listener Report (v2)               0           0          0
Other Unknown types                                       0
MLD v2 source required for SSM     2
MLD v2 mode not applicable for SSM 0

MLD Global Statistics
Bad Length                         0
Bad Checksum                       0
Bad Receive If                     0
Rx non-local                       0
Timed out                          0
Rejected Report                    0
Total Interfaces                   2
```

Copyright © 2018, Juniper Networks, Inc.

show msdp

Syntax

show msdp
<brief | detail>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<peer *peer-address*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display Multicast Source Discovery Protocol (MSDP) information.

Options

none—Display standard MSDP information for all routing instances.

brief | detail—(Optional) Display the specified level of output.

instance *instance-name*—(Optional) Display information for the specified instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

peer *peer-address*—(Optional) Display information about the specified peer only,

Required Privilege Level

view

Related Documentation

- show msdp source on page 1823
- show msdp source-active on page 1825
- show msdp statistics on page 1828

List of Sample Output

show msdp on page 1821
show msdp brief on page 1821
show msdp detail on page 1821

Output Fields

Table 114 on page 1820 describes the output fields for the **show msdp** command. Output fields are listed in the approximate order in which they appear.

Table 114: show msdp Output Fields

Field Name	Field Description	Level of Output
Peer address	IP address of the peer.	All levels
Local address	Local address of the peer.	All levels

 Copyright © 2018, Juniper Networks, Inc.

Table 114: show msdp Output Fields (continued)

Field Name	Field Description	Level of Output
State	Status of the MSDP connection: **Listen**, **Established**, or **Inactive**.	All levels
Last up/down	Time at which the most recent peer-state change occurred.	All levels
Peer-Group	Peer group name.	All levels
SA Count	Number of source-active cache entries advertised by each peer that were accepted, compared to the number that were received, in the format *number-accepted/number-received*.	All levels
Peer Connect Retries	Number of peer connection retries.	**detail**
State timer expires	Number of seconds before another message is sent to a peer.	**detail**
Peer Times out	Number of seconds to wait for a response from the peer before the peer is declared unavailable.	**detail**
SA accepted	Number of entries in the source-active cache accepted from the peer.	**detail**
SA received	Number of entries in the source-active cache received by the peer.	**detail**

Sample Output

show msdp

```
user@host> show msdp
Peer address    Local address    State       Last up/down Peer-Group SA Count
198.32.8.193    198.32.8.195     Established  5d 19:25:44 North23    120/150
198.32.8.194    198.32.8.195     Established  3d 19:27:27 North23    300/345
198.32.8.196    198.32.8.195     Established  5d 19:39:36 North23    10/13
198.32.8.197    198.32.8.195     Established  5d 19:32:27 North23    5/6
198.32.8.198    198.32.8.195     Established  3d 19:33:04 North23    2305/3000
```

show msdp brief

The output for the **show msdp brief** command is identical to that for the **show msdp** command. For sample output, see show msdp on page 1821.

show msdp detail

```
user@host> show msdp detail
Peer: 10.255.70.15
Local address: 10.255.70.19
State: Established
Peer Connect Retries: 0
State timer expires: 22
Peer Times out: 49
SA accepted: 0
SA received: 0
```

Copyright © 2018, Juniper Networks, Inc.

show msdp source

Syntax

show msdp source
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<*source-address*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display multicast sources learned from Multicast Source Discovery Protocol (MSDP).

Options

none—Display standard MSDP source information for all routing instances.

instance *instance-name*—(Optional) Display information for the specified instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

source-address—(Optional) IP address and optional prefix length. Display information for the specified source address only.

Required Privilege Level

view

Related Documentation

- show msdp on page 1820
- show msdp source-active on page 1825
- show msdp statistics on page 1828

List of Sample Output

show msdp source on page 1824

Output Fields

Table 115 on page 1824 describes the output fields for the **show msdp source** command. Output fields are listed in the approximate order in which they appear.

Table 115: show msdp source Output Fields

Field Name	Field Description
Source address	IP address of the source.
/Len	Length of the prefix for this IP address.
Type	Discovery method for this multicast source: • **Configured**—Source-active limit explicitly configured for this source. • **Dynamic**—Source-active limit established when this source was discovered.
Maximum	Source-active limit applied to this source.
Threshold	Source-active threshold applied to this source.
Exceeded	Number of source-active messages received from this source exceeding the established maximum.

Sample Output

show msdp source

```
user@host> show msdp source
Source address /Len   Type       Maximum   Threshold    Exceeded
0.0.0.0        /0     Configured       5       none        0
10.1.0.0       /16    Configured     500       none        0
10.1.1.1       /32    Configured   10000       none        0
10.1.1.2       /32    Dynamic       6936       none        0
10.1.5.5       /32    Dynamic        500       none          123
10.2.1.1       /32    Dynamic          2       none        0
```


Copyright © 2018, Juniper Networks, Inc.

show msdp source-active

Syntax

show msdp source-active
<brief | detail>
<group *group*>
<instance *instance-name*>
<local>
<logical-system (all | *logical-system-name*)>
<originator *originator*>
<peer *peer-address*>
<source *source-address*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display the Multicast Source Discovery Protocol (MSDP) source-active cache.

Options

none—Display standard MSDP source-active cache information for all routing instances.

brief | detail—(Optional) Display the specified level of output.

group *group*—(Optional) Display source-active cache information for the specified group.

instance *instance-name*—(Optional) Display information for the specified instance.

local—(Optional) Display all source-active caches originated by this router.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

originator *originator*—(Optional) Display information about the peer that originated the source-active cache entries.

peer *peer-address*—(Optional) Display the source-active cache of the specified peer.

source *source-address*—(Optional) Display the source-active cache of the specified source.

Required Privilege Level

view

Related Documentation

- show msdp on page 1820
- show msdp source on page 1823
- show msdp statistics on page 1828

List of Sample Output

show msdp source-active on page 1826
show msdp source-active brief on page 1827

show msdp source-active detail on page 1827
show msdp source-active source on page 1827

Output Fields

Table 116 on page 1826 describes the output fields for the **show msdp source-active** command. Output fields are listed in the approximate order in which they appear.

Table 116: show msdp source-active Output Fields

Field Name	Field Description
Global active source limit exceeded	Number of times all peers have exceeded configured active source limits.
Global active source limit maximum	Configured number of active source messages accepted by the device.
Global active source limit threshold	Configured threshold for applying random early discard (RED) to drop some but not all MSDP active source messages.
Global active source limit log-warning	Threshold at which a warning message is logged (percentage of the number of active source messages accepted by the device).
Global active source limit log interval	Time (in seconds) between consecutive log messages.
Group address	Multicast address of the group.
Source address	IP address of the source.
Peer address	IP address of the peer.
Originator	Router ID configured on the source of the rendezvous point (RP) that originated the message, or the loopback address when the router ID is not configured.
Flags	Flags: **Accept**, **Reject**, or **Filtered**.

Sample Output

show msdp source-active

```
user@host> show msdp source-active
Group address    Source address  Peer address    Originator      Flags
230.0.0.0        192.168.195.46  local           10.255.14.30    Accept
230.0.0.1        192.168.195.46  local           10.255.14.30    Accept
230.0.0.2        192.168.195.46  local           10.255.14.30    Accept
230.0.0.3        192.168.195.46  local           10.255.14.30    Accept
230.0.0.4        192.168.195.46  local           10.255.14.30    Accept
```


Copyright © 2018, Juniper Networks, Inc.

show msdp source-active brief

The output for the **show msdp source-active brief** command is identical to that for the **show msdp source-active** command. For sample output, see show msdp source-active on page 1826.

show msdp source-active detail

The output for the **show msdp source-active detail** command is identical to that for the **show msdp source-active** command. For sample output, see show msdp source-active on page 1826.

show msdp source-active source

```
user@host> show msdp source-active source 192.168.215.246
Global active source limit exceeded: 0
Global active source limit maximum: 25000
Global active source limit threshold: 24000
Global active source limit log-warning: 100
Global active source limit log interval: 0

Group address   Source address  Peer address    Originator      Flags
226.2.2.1       192.168.215.246 10.255.182.140  10.255.182.140  Accept
226.2.2.3       192.168.215.246 10.255.182.140  10.255.182.140  Accept
226.2.2.4       192.168.215.246 10.255.182.140  10.255.182.140  Accept
226.2.2.5       192.168.215.246 10.255.182.140  10.255.182.140  Accept
226.2.2.7       192.168.215.246 10.255.182.140  10.255.182.140  Accept
226.2.2.10      192.168.215.246 10.255.182.140  10.255.182.140  Accept
226.2.2.11      192.168.215.246 10.255.182.140  10.255.182.140  Accept
226.2.2.13      192.168.215.246 10.255.182.140  10.255.182.140  Accept
226.2.2.14      192.168.215.246 10.255.182.140  10.255.182.140  Accept
226.2.2.15      192.168.215.246 10.255.182.140  10.255.182.140  Accept
```

Copyright © 2018, Juniper Networks, Inc.

show msdp statistics

Syntax

show msdp statistics
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<peer *peer-address*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display statistics about Multicast Source Discovery Protocol (MSDP) peers.

Options

none—Display statistics about all MSDP peers for all routing instances.

instance *instance-name*—(Optional) Display statistics about a specific MSDP instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

peer *peer-address*—(Optional) Display statistics about a particular MSDP peer.

Required Privilege Level

view

Related Documentation

- clear msdp statistics on page 1740

List of Sample Output

show msdp statistics on page 1830
show msdp statistics peer on page 1830

Output Fields

Table 117 on page 1828 describes the output fields for the **show msdp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 117: show msdp statistics Output Fields

Field Name	Field Description
Global active source limit exceeded	Number of times all peers have exceeded configured active source limits.
Global active source limit maximum	Configured number of active source messages accepted by the device.
Global active source limit threshold	Configured threshold for applying random early discard (RED) to drop some but not all MSDP active source messages.
Global active source limit log-warning	Threshold at which a warning message is logged (percentage of the number of active source messages accepted by the device).

Copyright © 2018, Juniper Networks, Inc.

Table 117: show msdp statistics Output Fields (continued)

Field Name	Field Description
Global active source limit log interval	Time (in seconds) between consecutive log messages.
Peer	Address of peer.
Last State Change	How long ago the peer state changed.
Last message received from the peer	How long ago the last message was received from the peer.
RPF Failures	Number of reverse path forwarding (RPF) failures.
Remote Closes	Number of times the remote peer closed.
Peer Timeouts	Number of peer timeouts.
SA messages sent	Number of source-active messages sent.
SA messages received	Number of source-active messages received.
SA request messages sent	Number of source-active request messages sent.
SA request messages received	Number of source-active request messages received.
SA response messages sent	Number of source-active response messages sent.
SA response messages received	Number of source-active response messages received.
SA messages with zero Entry Count received	Entry Count is a field within SA message that defines how many source/group tuples are present in the SA message. The counter is incremented each time an SA with an Entry Count of zero is received.
Active source exceeded	Number of times this peer has exceeded configured source-active limits.
Active source Maximum	Configured number of active source messages accepted by this peer.
Active source threshold	Configured threshold on this peer for applying random early discard (RED) to drop some but not all MSDP active source messages.
Active source log-warning	Configured threshold on this peer at which a warning message is logged (percentage of the number of active source messages accepted by the device).
Active source log-interval	Time (in seconds) between consecutive log messages on this peer.

Copyright © 2018, Juniper Networks, Inc.

Table 117: show msdp statistics Output Fields (continued)

Field Name	Field Description
Keepalive messages sent	Number of keepalive messages sent.
Keepalive messages received	Number of keepalive messages received.
Unknown messages received	Number of unknown messages received.
Error messages received	Number of error messages received.

Sample Output

show msdp statistics

```
user@host> show msdp statistics
Global active source limit exceeded: 0
Global active source limit maximum: 10
Global active source limit threshold: 8
Global active source limit log-warning: 60
Global active source limit log interval: 60

Peer: 10.255.245.39
Last State Change: 11:54:49 (00:24:59)
Last message received from peer: 11:53:32 (00:26:16)
RPF Failures: 0
Remote Closes: 0
Peer Timeouts: 0
SA messages sent: 376
SA messages received: 459
SA messages with zero Entry Count received: 0
SA request messages sent: 0
SA request messages received: 0
SA response messages sent: 0
SA response messages received: 0
Active source exceeded: 0
Active source Maximum: 10
Active source threshold: 8
Active source log-warning: 60
Active source log-interval 120
Keepalive messages sent: 17
Keepalive messages received: 19
Unknown messages received: 0
Error messages received: 0
```

show msdp statistics peer

```
user@host> show msdp statistics peer 10.255.182.140
Peer: 10.255.182.140
    Last State Change: 8:19:23 (00:01:08)
    Last message received from peer: 8:20:05 (00:00:26)
    RPF Failures: 0
    Remote Closes: 0
    Peer Timeouts: 0
```

 Copyright © 2018, Juniper Networks, Inc.

```
SA messages sent: 17
SA messages received: 16
SA request messages sent: 0
SA request messages received: 0
SA response messages sent: 0
SA response messages received: 0
Active source exceeded: 20
Active source Maximum: 10
Active source threshold: 8
Active source log-warning: 60
Active source log-interval: 120
Keepalive messages sent: 0
Keepalive messages received: 0
Unknown messages received: 0
Error messages received: 0
```

Copyright © 2018, Juniper Networks, Inc.

show multicast backup-pe-groups

Syntax

show multicast backup-pe-groups
<address *pe-address*>
<group *group-name*>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced in Junos OS Release 9.0.

Description

Display backup PE router group information when ingress PE redundancy is configured. Ingress PE redundancy provides a backup resource when point-to-multipoint LSPs are configured for multicast distribution.

Options

none—Display standard information about all backup PE groups.

address ***pe-address***—(Optional) Display the groups that a PE address is associated with.

group ***group***—(Optional) Display the backup PE group information for a particular group.

instance ***instance-name***—(Optional) Display backup PE group information for a specific multicast instance.

logical-system (all | ***logical-system-name*****)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show multicast backup-pe-groups on page 1833

Output Fields

Table 118 on page 1832 describes the output fields for the **show multicast backup-pe-groups** command. Output fields are listed in the approximate order in which they appear.

Table 118: show multicast backup-pe-groups Output Fields

Field Name	Field Description
Backup PE Group	Group name.
Designated PE	Primary PE router. Address of the PE router that is currently forwarding traffic on the static route.
Transitions	Number of times that the designated PE router has transitioned from the most eligible PE router to a backup PE router and back again to the most eligible PE router.
Last Transition	Time of the most recent transition.
Local Address	Address of the local PE router.
Backup PE List	List of PE routers that are configured to be backups for the group.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show multicast backup-pe-groups

```
user@host> show multicast backup-pe-groups
Instance: master

Backup PE group: b1
    Designated PE:   10.255.165.7
    Transitions:     1
    Last Transition: 03:15:01
    Local Address:   10.255.165.7
    Backup PE List:
                     10.255.165.8

Backup PE group: b2
    Designated PE:   10.255.165.7
    Transitions:     2
    Last Transition: 02:58:20
    Local Address:   10.255.165.7
    Backup PE List:
                     10.255.165.9
                     10.255.165.8
```

Copyright © 2018, Juniper Networks, Inc.

show multicast flow-map

List of Syntax

Syntax on page 1834
Syntax (EX Series Switch and the QFX Series) on page 1834

Syntax

show multicast flow-map
<brief | detail>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)

show multicast flow-map
<brief | detail>

Release Information

Command introduced in Junos OS Release 8.2.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display configuration information about IP multicast flow maps.

Options

none—Display configuration information about IP multicast flow maps on all systems.

brief | detail—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show multicast flow-map on page 1835
show multicast flow-map detail on page 1835

Output Fields

Table 119 on page 1834 describes the output fields for the **show multicast flow-map** command. Output fields are listed in the approximate order in which they appear.

Table 119: show multicast flow-map Output Fields

Field Name	Field Description	Levels of Output
Name	Name of the flow map.	All levels
Policy	Name of the policy associated with the flow map.	All levels
Cache-timeout	Cache timeout value assigned to the flow map.	All levels
Bandwidth	Bandwidth setting associated with the flow map.	All levels
Adaptive	Whether or not adaptive mode is enabled for the flow map.	none

 Copyright © 2018, Juniper Networks, Inc.

Table 119: show multicast flow-map Output Fields (continued)

Field Name	Field Description	Levels of Output
Flow-map	Name of the flow map.	**detail**
Adaptive Bandwidth	Whether or not adaptive mode is enabled for the flow map.	**detail**
Redundant Sources	Redundant sources defined for the same destination group.	**detail**

Sample Output

show multicast flow-map

```
user@host> show multicast flow-map
Instance: master
Name            Policy                  Cache timeout        Bandwidth Adaptive
map2            policy2                         never          2000000 no
map1            policy1                    60 seconds          2000000 no
```

Sample Output

show multicast flow-map detail

```
user@host> show multicast flow-map detail
Instance: master
Flow-map: map1
    Policy:             policy1
    Cache Timeout:      600 seconds
    Bandwidth:          2000000
    Adaptive Bandwidth: yes
    Redundant Sources:  10.11.11.11
    Redundant Sources:  10.11.11.12
    Redundant Sources:  10.11.11.13
```

Copyright © 2018, Juniper Networks, Inc.

show multicast forwarding-cache statistics

Syntax

show multicast forwarding-cache statistics
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced in Junos OS Release 12.2.
Starting in Junos OS Release 16.1, output includes general and rendezvous-point tree (RPT) suppression states.

Description

Display IP multicast forwarding cache statistics.

Options

none—Display multicast forwarding cache statistics for all supported address families for all routing instances.

inet | inet6—(Optional) Display multicast forwarding cache statistics for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display multicast forwarding cache statistics for a specific routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- clear multicast forwarding-cache on page 1743
- *threshold*

List of Sample Output

show multicast forwarding cache statistics instance on page 1837
show multicast forwarding cache statistics instance (Forwarding-cache suppression is disabled) on page 1837

Output Fields

Table 120 on page 1836 describes the output fields for the **show multicast forwarding-cache statistics** command. Output fields are listed in the approximate order in which they appear.

Table 120: show multicast forwarding-cache statistics Output Fields

Field Name	Field Description
Instance	Name of the routing instance for which multicast forwarding cache statistics are displayed.
Family	Protocol family for which multicast forwarding cache statistics are displayed: **ALL**, **INET**, or **INET6**.

Copyright © 2018, Juniper Networks, Inc.

Table 120: show multicast forwarding-cache statistics Output Fields (continued)

Field Name	Field Description
General (or MVPN RPT) Suppression Active	Indicates whether suppression is configured.
General (or MVPN RPT) Entries Used	Number of currently used multicast forwarding cache entries.
General (or MVPN RPT) Suppress Threshold	Maximum number of multicast forwarding cache entries that can be added to the cache. When the number of entries reaches the configured threshold, the device suspends adding new multicast forwarding cache entries.
General (or MVPN RPT) Reuse Value	Number of multicast forwarding cache entries that must be reached before the device creates new multicast forwarding cache entries. When the total number of multicast forwarding cache entries is below the reuse value, the device resumes adding new multicast forwarding cache entries.

Sample Output

show multicast forwarding cache statistics instance

```
user@host> show multicast forwarding-cache statistic instance mvpn1 intet6
  Instance: mvpn1 Family: INET6
General Suppression Active                      Yes
General Entries Used                              0
General Suppress Threshold                      200
General Reuse Value                             200
MVPN RPT Suppression Active                     Yes
MVPN RPT Entries Used                             0
MVPN RPT Suppress Threshold                     200
MVPN RPT Reuse Value                            200
```

show multicast forwarding cache statistics instance (Forwarding-cache suppression is disabled)

```
user@host> show multicast forwarding-cache statistic instance mvpn1
Instance: mvpn1 Family: ALL
Forwarding-cache suppression disabled Not enabled by configuration
```

Copyright © 2018, Juniper Networks, Inc.

show multicast interface

List of Syntax

Syntax on page 1838
Syntax (EX Series Switch and the QFX Series) on page 1838

Syntax

show multicast interface
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)

show multicast interface

Release Information

Command introduced in Junos OS Release 8.3.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display bandwidth information about IP multicast interfaces.

Options

none—Display all interfaces that have multicast configured.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show multicast interface on page 1839

Output Fields

Table 121 on page 1838 describes the output fields for the **show multicast interface** command. Output fields are listed in the approximate order in which they appear.

Table 121: show multicast interface Output Fields

Field Name	Field Description
Interface	Name of the multicast interface.
Maximum bandwidth (bps)	Maximum bandwidth setting, in bits per second, for this interface.
Remaining bandwidth (bps)	Amount of bandwidth, in bits per second, remaining on the interface.

Copyright © 2018, Juniper Networks, Inc.

Table 121: show multicast interface Output Fields (continued)

Field Name	Field Description
Mapped bandwidth deduction (bps)	Amount of bandwidth, in bits per second, used by any flows that are mapped to the interface. NOTE: Adding the mapped bandwidth deduction value to the local bandwidth deduction value results in the total deduction value for the interface. This field does not appear in the output when the no QoS adjustment feature is disabled.
Local bandwidth deduction (bps)	Amount of bandwidth, in bits per second, used by any mapped flows that are traversing the interface. NOTE: Adding the mapped bandwidth deduction value to the local bandwidth deduction value results in the total deduction value for the interface. This field does not appear in the output when the no QoS adjustment feature is disabled.
Reverse OIF mapping	State of the reverse OIF mapping feature (**on** or **off**). NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.
Reverse OIF mapping no QoS adjustment	State of the no QoS adjustment feature (**on** or **off**) for interfaces that are using reverse OIF mapping. NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.
Leave timer	Amount of time a mapped interface remains active after the last mapping ends. NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.
No QoS adjustment	State (**on**) of the no QoS adjustment feature when this feature is enabled. NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.

Sample Output

show multicast interface

```
user@host> show multicast interface
Interface              Maximum bandwidth (bps) Remaining bandwidth (bps)
fe-0/0/3                              10000000                         0
fe-0/0/3.210                          10000000                  -2000000
fe-0/0/3.220                         100000000                 100000000
fe-0/0/3.230                          20000000                  18000000
fe-0/0/2.200                         100000000                 100000000
```

 Copyright © 2018, Juniper Networks, Inc.

show multicast mrinfo

Syntax

show multicast mrinfo
<*host*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display configuration information about IP multicast networks, including neighboring multicast router addresses.

Options

none—Display configuration information about all multicast networks.

host—(Optional) Display configuration information about a particular host. Replace *host* with a hostname or IP address.

Required Privilege Level

view

List of Sample Output

show multicast mrinfo on page 1842

Output Fields

Table 122 on page 1841 describes the output fields for the **show multicast mrinfo** command. Output fields are listed in the approximate order in which they appear.

Table 122: show multicast mrinfo Output Fields

Field Name	Field Description
source-address	Query address, hostname (DNS name or IP address of the source address), and multicast protocol version or the software version of another vendor.
ip-address-1—>ip-address-2	Queried router interface address and directly attached neighbor interface address, respectively.
(*name* or *ip-address*)	Name or IP address of neighbor.

[*metric/threshold/type/flags*]	Neighbor's multicast profile:
	• *metric*—Always has a value of **1**, because **mrinfo** queries the directly connected interfaces of a device.
	• *threshold*—Multicast threshold time-to-live (TTL). The range of values is **0** through **255**.
	• *type*—Multicast connection type: **pim** or **tunnel**.
	• *flags*—Flags for this route:
	  ◦ **querier**—Queried router is the designated router for the neighboring session.
	  ◦ **leaf**—Link is a leaf in the multicast network.
	  ◦ **down**—Link status indicator.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show multicast mrinfo

```
user@host> show multicast mrinfo 10.35.4.1
10.35.4.1 (10.35.4.1) [version 12.0]:
  192.168.195.166 -> 0.0.0.0 (local) [1/0/pim/querier/leaf]
  10.38.20.1 -> 0.0.0.0 (local) [1/0/pim/querier/leaf]
  10.47.1.1 -> 10.47.1.2 (10.47.1.2) [1/5/pim]
  0.0.0.0 -> 0.0.0.0 (local) [1/0/pim/down]
```


Copyright © 2018, Juniper Networks, Inc.

show multicast next-hops

List of Syntax

Syntax on page 1843
Syntax (EX Series Switch and the QFX Series) on page 1843

Syntax

show multicast next-hops
<brief | detail | terse>
<*identifier-number*>
<inet | inet6>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)

show multicast next-hops
<brief | detail>
<*identifier-number*>
<inet | inet6>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 option introduced in Junos OS Release 10.0 for EX Series switches.
detail option display of next-hop ID number introduced in Junos OS Release 11.1 for M Series and T Series routers and EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Support for bidirectional PIM added in Junos OS Release 12.1.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
terse option introduced in Junos OS Release 16.1 for the MX Series.

Description

Display the entries in the IP multicast next-hop table.

Options

none—Display standard information about all entries in the multicast next-hop table for all supported address families.

brief | detail | terse—(Optional) Display the specified level of output. Use **terse** to display the total number of outgoing interfaces (as opposed to listing them) When you include the **detail** option on M Series and T Series routers and EX Series switches, the downstream interface name includes the next-hop ID number in parentheses, in the form **fe-0/1/2.0-(1048574),** where **1048574** is the next-hop ID number.

Starting in Junos OS release 16.1, the **show multicast next-hops** statement shows the hierarchical next hops contained in the top-level next hop.

identifier-number—(Optional) Show a particular next hop by ID number. The range of values is **1** through **65,535**.

inet | inet6—(Optional) Display entries for IPv4 or IPv6 family addresses, respectively.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Copyright © 2018, Juniper Networks, Inc.

Required Privilege Level view

List of Sample Output
show multicast next-hops on page 1844
show multicast next-hops (Ingress Router, Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 1844
show multicast next-hops (Egress Router, Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 1845
show multicast next-hops (Bidirectional PIM) on page 1845
show multicast next-hops brief on page 1845
show multicast next-hops detail on page 1845

Output Fields Table 123 on page 1844 describes the output fields for the **show multicast next-hops** command. Output fields are listed in the approximate order in which they appear.

Table 123: show multicast next-hops Output Fields

Field Name	Field Description
Family	Protocol family (such as **INET**).
ID	Next-hop identifier of the prefix. The identifier is returned by the routing device's Packet Forwarding Engine.
Refcount	Number of cache entries that are using this next hop.
KRefcount	Kernel reference count for the next hop.
Downstream interface	Interface names associated with each multicast next-hop ID.
Incoming interface list	List of interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM.

Sample Output

show multicast next-hops

```
user@host> show multicast next-hops
Family: INET
ID       Refcount  KRefcount Downstream interface
262142          4          2 so-1/0/0.0
262143          2          1 mt-1/1/0.49152
262148          2          1 mt-1/1/0.32769
```

show multicast next-hops (Ingress Router, Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show multicast next-hops
Family: INET
ID           Refcount KRefcount Downstream interface Addr
1048580             2          1 1048576
(0x600dc04)         1          0 1048584
(0x600ea04)         1          0 (0x600e924)
```


Copyright © 2018, Juniper Networks, Inc.

```
1048583           2         1 1048579
(0x600e144)       1         0 1048587
(0x600e844)       1         0 (0x600e764)
1048582           2         1 1048578
(0x600df84)       1         0 1048586
(0x600e684)       1         0 (0x600e5a4)
1048581           2         1 1048577
(0x600ddc4)       1         0 1048585
(0x600ebc4)       1         0 (0x600eae4)
```

show multicast next-hops (Egress Router, Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show multicast next-hops
Family: INET
ID          Refcount KRefcount Downstream interface Addr
(0x600e844)       8         0 1048575
1048575          16         0 distributed-gmp
```

show multicast next-hops (Bidirectional PIM)

```
user@host> show multicast next-hops
Family: INET
ID          Refcount KRefcount Downstream interface
2097151           8         4 ge-0/0/1.0

Family: INET6
ID          Refcount KRefcount Downstream interface
2097157           2         1 ge-0/0/1.0

Family: Incoming interface list
ID          Refcount KRefcount Downstream interface
513               5         2 lo0.0
                              ge-0/0/1.0
514               5         2 lo0.0
                              ge-0/0/1.0
                              xe-4/1/0.0
515               3         1 lo0.0
                              ge-0/0/1.0
                              xe-4/1/0.0
544               1         0 lo0.0
                              xe-4/1/0.0
```

show multicast next-hops brief

The output for the **show multicast next-hops brief** command is identical to that for the **show multicast next-hops** command. For sample output, see show multicast next-hops on page 1844.

show multicast next-hops detail

```
user@host> show multicast next-hops detail
Family: INET
ID          Refcount KRefcount Downstream interface Addr
1048584           2         1 1048581
                              1048580
                              Flags 0x208 type 0x18 members 0/0/2/0/0
                              Address 0xb1841c4
1048591           3         2 787
                              747
```

```
                                  Flags 0x206 type 0x18 members 0/0/2/0/0
                                  Address 0xb1847f4
1048580            4          1 ge-1/1/9.0-(1048579)
                                  Flags 0x200 type 0x18 members 0/0/0/1/0
                                  Address 0xb184134
1048581            2          0 736
                                  765
                                  Flags 0x3 type 0x18 members 0/0/2/0/0
                                  Address 0xb183dd4
1048585           18          0 787
                                  747
                                  Flags 0x203 type 0x18 members 0/0/2/0/0
                                  Address 0xb184404

Family: INET6
ID          Refcount KRefcount Downstream interface Addr
1048586            4          2 1048585
                                  1048583
                                  Flags 0x20c type 0x19 members 0/0/2/0/0
                                  Address 0xb1842e4
1048583           14          4 ge-1/1/9.0-(1048582)
                                  Flags 0x200 type 0x19 members 0/0/0/1/0
                                  Address 0xb183ef4
1048592            4          2 1048583
                                  1048591
                                  Flags 0x20c type 0x19 members 0/0/2/0/0
                                  Address 0xb184644
```

 Copyright © 2018, Juniper Networks, Inc.

show multicast pim-to-igmp-proxy

List of Syntax

Syntax on page 1847
Syntax (EX Series Switch and the QFX Series) on page 1847

Syntax

show multicast pim-to-igmp-proxy
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)

show multicast pim-to-igmp-proxy
<instance *instance-name*>

Release Information

Command introduced in Junos OS Release 9.6.
Command introduced in Junos OS Release 9.6 for EX Series switches.
instance option introduced in Junos OS Release 10.3.
instance option introduced in Junos OS Release 10.3 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display configuration information about PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy.

Options

none—Display configuration information about PIM-to-IGMP message translation for all routing instances.

instance *instance-name*—(Optional) Display configuration information about PIM-to-IGMP message translation for a specific multicast instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- *Configuring PIM-to-IGMP and PIM-to-MLD Message Translation*

List of Sample Output

show multicast pim-to-igmp-proxy on page 1848
show multicast pim-to-igmp-proxy instance on page 1848

Output Fields

Table 124 on page 1848 describes the output fields for the **show multicast pim-to-igmp-proxy** command. Output fields are listed in the order in which they appear.

Table 124: show multicast pim-to-igmp-proxy Output Fields

Field Name	Field Description
Instance	Routing instance. Default instance is **master** (inet.0 routing table).
Proxy state	State of PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy, on the configured upstream interfaces: **enabled** or **disabled**.
interface-name	Name of upstream interface (no more than two allowed) on which PIM-to-IGMP message translation is configured.

Sample Output

show multicast pim-to-igmp-proxy

```
user@host> show multicast pim-to-igmp-proxy
Instance: master Proxy state: enabled
ge-0/1/0.1
ge-0/1/0.2
```

show multicast pim-to-igmp-proxy instance

```
user@host> show multicast pim-to-igmp-proxy instance VPN-A
Instance: VPN-A Proxy state: enabled
ge-0/1/0.1
```

 Copyright © 2018, Juniper Networks, Inc.

show multicast pim-to-mld-proxy

List of Syntax

Syntax on page 1849
Syntax (EX Series Switch and the QFX Series) on page 1849

Syntax

show multicast pim-to-mld-proxy
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)

show multicast pim-to-mld-proxy
<instance *instance-name*>

Release Information

Command introduced in Junos OS Release 9.6.
Command introduced in Junos OS Release 9.6 for EX Series switches.
instance option introduced in Junos OS Release 10.3.
instance option introduced in Junos OS Release 10.3 for EX Series switches.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

Display configuration information about PIM-to-MLD message translation, also known as PIM-to-MLD proxy.

Options

none—Display configuration information about PIM-to-MLD message translation for all routing instances.

instance *instance-name*—(Optional) Display configuration information about PIM-to-MLD message translation for a specific multicast instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show multicast pim-to-mld-proxy on page 1850
show multicast pim-to-mld-proxy instance on page 1850

Output Fields

Table 125 on page 1849 describes the output fields for the **show multicast pim-to-mld-proxy** command. Output fields are listed in the order in which they appear.

Table 125: show multicast pim-to-mld-proxy Output Fields

Field Name	Field Description
Proxy state	State of PIM-to-MLD message translation, also known as PIM-to-MLD proxy, on the configured upstream interfaces: **enabled** or **disabled**.

Copyright © 2018, Juniper Networks, Inc.

Table 125: show multicast pim-to-mld-proxy Output Fields (continued)

Field Name	Field Description
interface-name	Name of upstream interface (no more than two allowed) on which PIM-to-MLD message translation is configured.

Sample Output

show multicast pim-to-mld-proxy

```
user@host> show multicast pim-to-mld-proxy
Instance: master Proxy state: enabled
ge-0/5/0.1
ge-0/5/0.2
```

show multicast pim-to-mld-proxy instance

```
user@host> show multicast pim-to-mld-proxy instance VPN-A
Instance: VPN-A Proxy state: enabled
ge-0/5/0.1
```


Copyright © 2018, Juniper Networks, Inc.

show multicast route

List of Syntax

Syntax on page 1851
Syntax (EX Series Switch and the QFX Series) on page 1851

Syntax

show multicast route
<brief | detail | extensive | summary>
<active | all | inactive>
<group *group*>
<inet | inet6>
<instance *instance name*>
<logical-system (all | *logical-system-name*)>
<oif-count>
<*regular-expression*>
<source-prefix *source-prefix*>

Syntax (EX Series Switch and the QFX Series)

show multicast route
<brief | detail | extensive | summary>
<active | all | inactive>
<group *group*>
<inet | inet6>
<instance *instance name*>
<*regular-expression*>
<source-prefix *source-prefix*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Support for bidirectional PIM added in Junos OS Release 12.1.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
oif-count option introduced in Junos OS Release 16.1 for the MX Series.
xxxSupport for PIM NSR support for VXLAN added in Junos OS Release 16.2.

Description

Display the entries in the IP multicast forwarding table. You can display similar information with the **show route table inet.1** command.

NOTE: On all SRX Series devices, when a multicast route is not available, pending sessions are not torn down, and subsequent packets are queued. If no multicast route resolve comes back, then the traffic flow has to wait for the pending session to timed out. Then packets can trigger new pending session create and route resolve.

Options

none—Display standard information about all entries in the multicast forwarding table for all routing instances.

brief | detail | extensive | summary—(Optional) Display the specified level of output.

active | all | inactive—(Optional) Display all active entries, all entries, or all inactive entries, respectively, in the multicast forwarding table.

group ***group***—(Optional) Display the cache entries for a particular group.

inet | inet6—(Optional) Display multicast forwarding table entries for IPv4 or IPv6 family addresses, respectively.

instance ***instance-name***—(Optional) Display entries in the multicast forwarding table for a specific multicast instance.

logical-system (all | ***logical-system-name*****)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

oif-count —(Optional) Display a count of outgoing interfaces rather than listing them.

regular-expression—(Optional) Display information about the multicast forwarding table entries that match a UNIX OS-style regular expression.

source-prefix ***source-prefix***—(Optional) Display the cache entries for a particular source prefix.

Required Privilege Level

view

Related Documentation

- *Example: Configuring Multicast-Only Fast Reroute in a PIM Domain*

List of Sample Output

Output Fields

Table 126 on page 1853 describes the output fields for the **show multicast route** command. Output fields are listed in the approximate order in which they appear.

 Copyright © 2018, Juniper Networks, Inc.

Table 126: show multicast route Output Fields

Field Name	Field Description	Level of Output
family	IPv4 address family (**INET**) or IPv6 address family (**INET6**).	All levels
Group	Group address. For any-source multicast routes, for example for bidirectional PIM, the group address includes the prefix length.	All levels
Source	Prefix and length of the source as it is in the multicast forwarding table.	All levels
Incoming interface list	List of interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM.	All levels
Upstream interface	Name of the interface on which the packet with this source prefix is expected to arrive.	All levels
Upstream rpf interface list	When multicast-only fast reroute (MoFRR) is enabled, a PIM router propagates join messages on two upstream RPF interfaces to receive multicast traffic on both links for the same join request.	All levels
Downstream interface list	List of interface names to which the packet with this source prefix is forwarded. **distributed-gmp**— Added in Junos OS Release 17.4R1 to indicate that line cards with *distributed* IGMP interfaces are receiving multicast traffic for a given (s,g).	All levels
Number of outgoing interfaces	Total number of outgoing interfaces for each (S,G) entry.	**extensive**
Session description	Name of the multicast session.	**detail extensive**
Statistics	Rate at which packets are being forwarded for this source and group entry (in Kbps and pps), and number of packets that have been forwarded to this prefix. If one or more of the kilobits per second packet forwarding statistic queries fails or times out, the statistics field displays **Forwarding statistics are not available**. NOTE: On QFX Series switches and OCX Series switches, this field does not report valid statistics.	**detail extensive**
Next-hop ID	Next-hop identifier of the prefix. The identifier is returned by the routing device's Packet Forwarding Engine and is also displayed in the output of the **show multicast nexthops** command.	**detail extensive**
Incoming interface list ID	For bidirectional PIM, incoming interface list identifier. Identifiers for interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM.	**detail extensive**

Table 126: show multicast route Output Fields (continued)

Field Name	Field Description	Level of Output
Upstream protocol	The protocol that maintains the active multicast forwarding route for this group or source. When the **show multicast route extensive** command is used with the **display-origin-protocol** option, the field name is only **Protocol** and not **Upstream Protocol**. However, this field also displays the protocol that installed the active route.	**detail extensive**
Route type	Type of multicast route. Values can be (S,G) or (*,G).	**summary**
Route state	Whether the group is **Active** or **Inactive**.	**summary extensive**
Route count	Number of multicast routes.	**summary**
Forwarding state	Whether the prefix is pruned or forwarding.	**extensive**
Cache lifetime/timeout	Number of seconds until the prefix is removed from the multicast forwarding table. A value of **never** indicates a permanent forwarding entry. A value of **forever** indicates routes that do not have keepalive times.	**extensive**
Wrong incoming interface notifications	Number of times that the upstream interface was not available.	**extensive**
Uptime	Time since the creation of a multicast route.	**extensive**

Sample Output

Starting in Junos OS Release16.1, **show multicast route** displays the top-level hierarchical next hop.

show multicast route

```
user@host> show multicast route
Family: INET

Group: 233.252.0.0
    Source: 10.255.14.144/32
    Upstream interface: local
    Downstream interface list:
        so-1/0/0.0

Group: 233.252.0.1
    Source: 10.255.14.144/32
    Upstream interface: local
    Downstream interface list:
        so-1/0/0.0

Group: 233.252.0.1
    Source: 10.255.70.15/32
    Upstream interface: so-1/0/0.0
    Downstream interface list:
```

 Copyright © 2018, Juniper Networks, Inc.

```
        mt-1/1/0.1081344

Family: INET6
```

show multicast route (Bidirectional PIM)

```
user@host> show multicast route
Family: INET

Group: 233.252.0.1/24
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0
    Downstream interface list:
        ge-0/0/1.0

Group: 233.252.0.3/24
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0 xe-4/1/0.0
    Downstream interface list:
        ge-0/0/1.0

Group: 233.252.0.11/24
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0
    Downstream interface list:
        ge-0/0/1.0

Group: 233.252.0.13/24
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0 xe-4/1/0.0
    Downstream interface list:
        ge-0/0/1.0
Family: INET6
```

show multicast route brief

The output for the **show multicast route brief** command is identical to that for the **show multicast route** command. For sample output, see show multicast route on page 1854 or show multicast route (Bidirectional PIM) on page 1855.

show multicast route summary

```
user@host>show multicast route summary
Instance: master Family: INET

Route type   Route state    Route count
(S,G)             Active            2
(S,G)             Inactive          3

Instance: master Family: INET6
```

show multicast route detail

```
user@host> show multicast route detail
```

Copyright © 2018, Juniper Networks, Inc.

```
Family: INET

Group: 233.252.0.0
    Source: 10.255.14.144/32
    Upstream interface: local
    Downstream interface list:
        so-1/0/0.0
    Session description: Unknown
    Statistics: 8 kBps, 100 pps, 45272 packets
    Next-hop ID: 262142
    Upstream protocol: PIM

Group: 233.252.0.1
    Source: 10.255.14.144/32
    Upstream interface: local
    Downstream interface list:
        so-1/0/0.0
    Session description: Administratively Scoped
    Statistics: 0 kBps, 0 pps, 13404 packets
    Next-hop ID: 262142
    Upstream protocol: PIM

Group: 233.252.0.1
    Source: 10.255.70.15/32
    Upstream interface: so-1/0/0.0
    Downstream interface list:
        mt-1/1/0.1081344
    Session description: Administratively Scoped
    Statistics: 46 kBps, 1000 pps, 921077 packets

    Next-hop ID: 262143
    Upstream protocol: PIM

Family: INET6
```

show multicast route extensive (Bidirectional PIM)

```
user@host> show multicast route extensive
Family: INET

Group: 233.252.0.1/24
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0
    Downstream interface list:
        ge-0/0/1.0
    Number of outgoing interfaces: 1
    Session description: NOB Cross media facilities
    Statistics: 0 kBps, 0 pps, 0 packets
    Next-hop ID: 2097153
    Incoming interface list ID: 585
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0

Group: 233.252.0.3/24
    Source: *
    Incoming interface list:
        lo0.0 ge-0/0/1.0 xe-4/1/0.0
```


Copyright © 2018, Juniper Networks, Inc.

```
    Downstream interface list:
        ge-0/0/1.0
    Number of outgoing interfaces: 1
    Session description: NOB Cross media facilities
    Statistics: 0 kBps, 0 pps, 0 packets
    Next-hop ID: 2097153
    Incoming interface list ID: 589
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0

Family: INET6
```

show multicast route extensive (Ingress Router, Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show multicast route extensive
 Family: INET

Group: 226.0.0.1
    Source: 200.1.0.2/32
    Upstream interface: xe-3/0/0.0
    Downstream interface list:
        ge-0/1/9.0
    Number of outgoing interfaces: 1
    Session description: Unknown
    Statistics: 78 kBps, 1000 pps, 34789 packets
    Next-hop ID: 1048582
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: 360 seconds
    Wrong incoming interface notifications: 0
    Uptime: 00:00:35

Group: 226.0.0.2
    Source: 200.1.0.2/32
    Upstream interface: xe-3/0/0.0
    Downstream interface list:
        ge-0/1/9.0
    Number of outgoing interfaces: 1
    Session description: Unknown
    Statistics: 78 kBps, 1000 pps, 34788 packets
    Next-hop ID: 1048583
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: 360 seconds
    Wrong incoming interface notifications: 0
    Uptime: 00:00:35

Group: 226.0.0.3
    Source: 200.1.0.2/32
    Upstream interface: xe-3/0/0.0
    Downstream interface list:
        ge-0/1/9.0
    Number of outgoing interfaces: 1
    Session description: Unknown
    Statistics: 78 kBps, 1000 pps, 34786 packets
    Next-hop ID: 1048580
```

Copyright © 2018, Juniper Networks, Inc.

```
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: 360 seconds
    Wrong incoming interface notifications: 0
    Uptime: 00:00:35

Group: 226.0.0.4
    Source: 200.1.0.2/32
    Upstream interface: xe-3/0/0.0
    Downstream interface list:
        ge-0/1/9.0
    Number of outgoing interfaces: 1
    Session description: Unknown
    Statistics: 78 kBps, 1000 pps, 34787 packets
    Next-hop ID: 1048581
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: 360 seconds
    Wrong incoming interface notifications: 0
    Uptime: 00:00:35

Instance: master Family: INET6
```

show multicast route instance <instance-name> extensive

```
user@host> show multicast route instance mvpn extensive
Family: INET
roup: 233.252.0.10
    Source: 10.0.0.2/32
    Upstream interface: xe-0/0/0.102
    Downstream interface list:
        xe-10/3/0.0 xe-0/3/0.0 xe-0/0/0.106 xe-0/0/0.105
        xe-0/0/0.103 xe-0/0/0.104 xe-0/0/0.107 xe-0/0/0.108
    Session description: Administratively Scoped
    Statistics: 256 kBps, 3998 pps, 670150 packets
    Next-hop ID: 1048579
    Upstream protocol: MVPN
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 58
    Uptime: 00:00:04

Instance: master Family: INET

Group: 225.0.0.1
    Source: 101.0.0.2/32
    Upstream interface: ge-2/2/0.101
    Downstream interface list:
        distributed-gmp
    Number of outgoing interfaces: 1
    Session description: Unknown
    Statistics: 105 kBps, 2500 pps, 4153361 packets
    Next-hop ID: 1048575
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: 360 seconds
    Wrong incoming interface notifications: 0
```

 Copyright © 2018, Juniper Networks, Inc.

```
    Uptime: 00:31:46

Group: 225.0.0.1
    Source: 101.0.0.3/32
    Upstream interface: ge-2/2/0.101
    Downstream interface list:
        distributed-gmp
    Number of outgoing interfaces: 1
    Session description: Unknown
    Statistics: 105 kBps, 2500 pps, 4153289 packets
    Next-hop ID: 1048575
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: 360 seconds
    Wrong incoming interface notifications: 0
    Uptime: 00:31:46
```

show multicast route extensive instance <instance-name>

`user@host>` **show multicast route extensive instance VPN-A**

The double asterisks (**) indicate new output, specific to **min-rate** and **revert-delay** settings made under the **[edit routing-instances** *routing-instance-name* protocols mvpn *hot-root-standby*] hierarchy.

```
Instance: VPNA Family: INET

      Group: 227.1.1.1
          Source: 18.1.1.2/32
          Upstream rpf interface list:
              vt-2/0/10.1000 (P)
                Session Id: 0x156 Session Status: Up
                Min-rate: 10000 kbps Weight: 1
                  Sender Id: Label 299808
              vt-2/0/10.1000 (B)
                Session Id: 0x155 Session Status: Up
                Min-rate: 10000 kbps Weight: 65533
                Sender Id: Label 299824
          Downstream interface list:
              lt-2/0/10.0
          Number of outgoing interfaces: 1
          Session description: Unknown
          Statistics: 8258 kBps, 100707 pps, 513032034 packets
          RPF Next-hop ID: 803
          Next-hop ID: 1048580
          Upstream protocol: MVPN
          Route state: Active
          Forwarding state: Forwarding
          Cache lifetime/timeout: forever
          Wrong incoming interface notifications: 0
          Uptime: 01:24:55
```

show multicast route extensive (PIM NSR support for VXLAN on master Routing Engine)

`user@host>` **show multicast route extensive**

```
Instance: master Family: INET

Group: 233.252.0.1
    Source: 10.3.3.3/32
```

```
    Upstream interface: ge-3/1/2.0
    Downstream interface list:
        -(593)
    Number of outgoing interfaces: 1
    Session description: Organizational Local Scope
    Statistics: 0 kBps, 0 pps, 27 packets
    Next-hop ID: 1048576
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding (Forwarding state is set as 'Forwarding' in
master RE.)
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0
    Uptime: 00:06:38

Group: 233.252.0.1
    Source: 10.2.1.4/32
    Upstream interface: local
    Downstream interface list:
        ge-3/1/2.0
    Number of outgoing interfaces: 1
    Session description: Organizational Local Scope
    Statistics: 0 kBps, 0 pps, 86 packets
    Next-hop ID: 1048575
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding (Forwarding state is set as 'Forwarding' in
master RE.)
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0
    Uptime: 00:07:45

Instance: master Family: INET6
```

show multicast route extensive (PIM NSR support for VXLAN on backup Routing Engine)

```
user@host> show multicast route extensive
Instance: master Family: INET

Group: 233.252.0.1
    Source: 10.3.3.3/32
    Upstream interface: ge-3/1/2.0
    Number of outgoing interfaces: 0
    Session description: Organizational Local Scope
    Forwarding statistics are not available
    Next-hop ID: 0
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Pruned (Forwarding state is set as 'Pruned' in backup RE.)

    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0
    Uptime: 00:06:46

Group: 233.252.0.1
    Source: 10.2.1.4/32
    Upstream interface: local
    Number of outgoing interfaces: 0
    Session description: Organizational Local Scope
    Forwarding statistics are not available
    Next-hop ID: 0
```


Copyright © 2018, Juniper Networks, Inc.

```
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Pruned  (Forwarding state is set as 'Pruned' in backup RE.)

    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0
    Uptime: 00:07:54

Instance: master Family: INET6
```

show multicast route extensive (PIM NSR support for VXLAN on backup Routing Engine)

```
user@host> show multicast route extensive
Instance: master Family: INET

Group: 233.252.0.1
    Source: 10.3.3.3/32
    Upstream interface: ge-3/1/2.0
    Downstream interface list:
        -(593)
    Number of outgoing interfaces: 1
    Session description: Organisational Local Scope
    Statistics: 0 kBps, 0 pps, 0 packets
    Next-hop ID: 1048576
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding (Forwarding state is set as 'Forwarding' in
backup RE.)
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0
    Uptime: 00:06:38

Group: 233.252.0.1
    Source: 10.2.1.4/32
    Upstream interface: local
    Downstream interface list:
        ge-3/1/2.0
    Number of outgoing interfaces: 1
    Session description: Organisational Local Scope
    Statistics: 0 kBps, 0 pps, 0 packets
    Next-hop ID: 1048575
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding (Forwarding state is set as 'Forwarding' in
backup RE.)
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0
    Uptime: 00:07:45

Instance: master Family: INET6
```

Copyright © 2018, Juniper Networks, Inc.

show multicast rpf

List of Syntax

Syntax on page 1862
Syntax (EX Series Switch and the QFX Series) on page 1862

Syntax

show multicast rpf
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<*prefix*>
<summary>

Syntax (EX Series Switch and the QFX Series)

show multicast rpf
<inet | inet6>
<instance *instance-name*>
<*prefix*>
<summary>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display information about multicast reverse-path-forwarding (RPF) calculations.

Options

none—Display RPF calculation information for all supported address families.

inet | inet6—(Optional) Display the RPF calculation information for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display information about multicast RPF calculations for a specific multicast instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

prefix—(Optional) Display the RPF calculation information for the specified prefix.

summary—(Optional) Display a summary of all multicast RPF information.

Required Privilege Level

view

List of Sample Output

show multicast rpf on page 1863
show multicast rpf inet6 on page 1864
show multicast rpf prefix on page 1865
show multicast rpf summary on page 1865

Copyright © 2018, Juniper Networks, Inc.

Output Fields Table 127 on page 1863 describes the output fields for the **show multicast rpf** command. Output fields are listed in the approximate order in which they appear.

Table 127: show multicast rpf Output Fields

Field Name	Field Description
Instance	Name of the routing instance. (Displayed when multicast is configured within a routing instance.)
Source prefix	Prefix and length of the source as it exists in the multicast forwarding table.
Protocol	How the route was learned.
Interface	Upstream RPF interface. NOTE: The displayed interface information does not apply to bidirectional PIM RP addresses. This is because the **show multicast rpf** command does not take into account equal-cost paths or the designated forwarder. For accurate upstream RPF interface information, always use the **show pim join extensive** command when bidirectional PIM is configured.
Neighbor	Upstream RPF neighbor. NOTE: The displayed neighbor information does not apply to bidirectional PIM. This is because the **show multicast rpf** command does not take into account equal-cost paths or the designated forwarder. For accurate upstream RPF neighbor information, always use the **show pim join extensive** command when bidirectional PIM is configured.

Sample Output

show multicast rpf

```
user@host> show multicast rpf

Multicast RPF table: inet.0, 12 entries

0.0.0.0/0
    Protocol: Static

10.255.14.132/32
    Protocol: Direct
    Interface: lo0.0

10.255.245.91/32
    Protocol: IS-IS
    Interface: so-1/1/1.0
    Neighbor: 192.168.195.21

172.16.0.1/32
Inactive172.16.0.0/12
Protocol: Static
Interface: fxp0.0
```

Copyright © 2018, Juniper Networks, Inc.

```
Neighbor: 192.168.14.254

192.168.0.0/16
Protocol: Static
Interface: fxp0.0
Neighbor: 192.168.14.254

192.168.14.0/24
Protocol: Direct
Interface: fxp0.0

192.168.14.132/32
Protocol: Local

192.168.195.20/30
Protocol: Direct
Interface: so-1/1/1.0

192.168.195.22/32
Protocol: Local

192.168.195.36/30
Protocol: IS-IS
Interface: so-1/1/1.0
Neighbor: 192.168.195.21
```

show multicast rpf inet6

user@host> **show multicast rpf inet6**

```
Multicast RPF table: inet6.0, 12 entries

::10.255.14.132/128
    Protocol: Direct
    Interface: lo0.0

::10.255.245.91/128
Protocol: IS-IS
Interface: so-1/1/1.0
Neighbor: 2001:db8::2a0:a5ff:fe28:2e8c

::192.168.195.20/126
Protocol: Direct
Interface: so-1/1/1.0

::192.168.195.22/128
Protocol: Local

::192.168.195.36/126
Protocol: IS-IS
Interface: so-1/1/1.0
Neighbor: 2001:db8::2a0:a5ff:fe28:2e8c

::192.168.195.76/126
Protocol: Direct
Interface: fe-2/2/0.0

::192.168.195.77/128
Protocol: Local
```

 Copyright © 2018, Juniper Networks, Inc.

```
2001:db8::/64
Protocol: Direct
Interface: so-1/1/1.0

2001:db8::290:69ff:fe0c:993a/128
Protocol: Local

2001:db8::2a0:a5ff:fe12:84f/128
Protocol: Direct
Interface: lo0.0

2001:db8::2/128
Protocol: PIM

2001:db8::d/128
Protocol: PIM
```

show multicast rpf prefix

```
user@host> show multicast rpf 2001:db8::/16

Multicast RPF table: inet6.0, 13 entries

2001:db8::2/128
    Protocol: PIM

2001:db8::d/128
    Protocol: PIM

...
```

show multicast rpf summary

```
user@host> show multicast rpf summary

Multicast RPF table: inet.0, 16 entries
Multicast RPF table: inet6.0, 12 entries
```

Copyright © 2018, Juniper Networks, Inc.

show multicast scope

List of Syntax

Syntax on page 1866
Syntax (EX Series Switch and the QFX Series) on page 1866

Syntax

show multicast scope
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)

show multicast scope
<inet | inet6>
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display administratively scoped IP multicast information.

Options

none—Display standard information about administratively scoped multicast information for all supported address families in all routing instances.

inet | inet6—(Optional) Display scoped multicast information for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display administratively scoped information for a specific multicast instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show multicast scope on page 1867
show multicast scope inet on page 1867
show multicast scope inet6 on page 1867

Output Fields

Table 128 on page 1866 describes the output fields for the **show multicast scope** command. Output fields are listed in the approximate order in which they appear.

Table 128: show multicast scope Output Fields

Field Name	Field Description
Scope name	Name of the multicast scope.

Copyright © 2018, Juniper Networks, Inc.

Table 128: show multicast scope Output Fields (continued)

Field Name	Field Description
Group Prefix	Range of multicast groups that are scoped.
Interface	Interface that is the boundary of the administrative scope.
Resolve Rejects	Number of kernel resolve rejects.

Sample Output

show multicast scope

```
user@host> show multicast scope
                                                                     Resolve
Scope name              Group Prefix       Interface                 Rejects
233-net                 233.252.0.0/16     fe-0/0/0.1                      0
local                   233.252.0.1/16     fe-0/0/0.1                      0
local                   2001:db8::/16          fe-0/0/0.1                      0
larry                   2001:db8::1234/128     fe-0/0/0.1                      0
```

show multicast scope inet

```
user@host> show multicast scope inet
                                                                     Resolve
Scope name              Group Prefix       Interface                 Rejects
233-net                 233.252.0.0/16     fe-0/0/0.1                      0
local                   233.252.0.0/16     fe-0/0/0.1                      0
```

show multicast scope inet6

```
user@host> show multicast scope inet6
                                                                     Resolve
Scope name              Group Prefix       Interface                 Rejects
local                   2001:db8::/16          fe-0/0/0.1                      0
larry                   2001:db8::1234/128     fe-0/0/0.1                      0
```

show multicast snooping next-hops

Syntax

show multicast snooping next-hops
<brief | detail>
<identifier *next-hop-ID*>
<inet>
<inet6>
<logical-system *logical-system-name*>

Release Information

Command introduced in Junos OS Release 11.2.

Description

Display information about the IP multicast snooping next-hops.

Options

brief | detail—(Optional) Display the specified level of output.

inet—(Optional) Display information for IPv4 multicast next hops only. If a family is not specified, both IPv4 and IPv6 results will be shown.

inet6—(Optional) Display information for IPv6 multicast next hops only. If a family is not specified, both IPv4 and IPv6 results will be shown.

logical-system ***logical-system-name***—(Optional) Display information about a particular logical system, or type 'all'.

Required Privilege Level

view

List of Sample Output

show multicast snooping next-hops on page 1870
show multicast snooping next-hops (IGMP snooping enabled on a VPLS) on page 1870

Output Fields

Table 129 on page 1868 describes the output fields for the **show multicast snooping next-hops** command. Output fields are listed in the approximate order in which they appear.

Table 129: show multicast snooping next-hops Output Fields

Field Name	Field Description
Family	Protocol family for which multicast snooping next hops are displayed: **INET** or **INET6**.
Refcount	Number of cache entries that are using this next hop.
KRefcount	Kernel reference count for the next hop.
Downstream interface	Interface names associated with each multicast next-hop ID.

Copyright © 2018, Juniper Networks, Inc.

Table 129: show multicast snooping next-hops Output Fields (continued)

Field Name	Field Description
Nexthop Id	Identifier for the next-hop. NOTE: To see the next-hop ID for a given PE mesh group, **igmp-snooping** must be enabled for the relevant VPLS routing instance. (Junos OS creates a default CE and VE mesh groups for each VPLS routing instance. The next hop of the VE mesh group is the set of VE mesh-group interfaces of the remaining PEs in the same VPLS routing instance.)

Sample Output

show multicast snooping next-hops

```
user@host> show multicast snooping next-hops
Family: INET
ID          Refcount KRefcount Downstream interface Nexthop Id
1048574            4         1 ge-0/1/0.1000
                               ge-0/1/2.1000
                               ge-0/1/3.1000

1048574            4         1 ge-0/1/0.1000-(2000)
                               1048575
                               1048576

1048575            2         0 ge-0/1/2.1000-(2001)
                               ge-0/1/3.1000-(2002)

1048576            2         0 lsi.1048578-(2003)
                               lsi.1048579-(2004)
```

show multicast snooping next-hops (IGMP snooping enabled on a VPLS)

In thIS example, ID 1048585 is the VE next-hop ID created for the VE next hop that is holding VE interfaces for the routing instance. It only appears if igmp snooping is enabled on the VPLS.

```
user@host> show multicast snooping next-hops
Family: INET
ID          Refcount KRefcount Downstream interface Addr
1048588            2         1 1048585
1048589            2         1 1048585
                               ge-0/0/5.100
0                  2         0 ge-0/0/0.100
                               ge-0/0/1.100
1048583            2         1 local
1048587            2         1 local
                               1048585
1048586            4         2 local
                               1048585
                               ge-0/0/5.100
1048584            2         1 local
                               ge-0/0/5.100
1048582            6         2 ge-0/0/5.100
0                  2         0 ge-0/0/0.200
                               ge-0/0/2.200
0                  2         0 ge-0/0/0.300
                               ge-0/0/2.300
0                  1         0 vt-0/0/10.17825792
                               vt-0/0/10.17825793
0                  1         0 vt-0/0/10.1048576
                               vt-0/0/10.1048578
1048585            5         0 vt-0/0/10.1048577
                               vt-0/0/10.1048579
0                  1         0 vt-0/0/10.34603008
                               vt-0/0/10.34603009
```

 Copyright © 2018, Juniper Networks, Inc.

show multicast sessions

List of Syntax

Syntax on page 1871
Syntax (EX Series Switch and the QFX Series) on page 1871

Syntax

show multicast sessions
<brief | detail | extensive>
<logical-system (all | *logical-system-name*)>
<*regular-expression*>

Syntax (EX Series Switch and the QFX Series)

show multicast sessions
<brief | detail | extensive>
<*regular-expression*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display information about announced IP multicast sessions.

NOTE: On all SRX Series devices, only 100 packets can be queued during pending (S, G) route. However, when multiple multicast sessions enter the route resolve process at the same time, buffer resources are not sufficient to queue 100 packets for each session.

Options

none— Display standard information about all multicast sessions for all routing instances.

brief | detail | extensive—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

regular-expression—(Optional) Display information about announced sessions that match a UNIX-style regular expression.

Required Privilege Level

view

List of Sample Output

show multicast sessions on page 1873
show multicast sessions regular-expression detail on page 1873

Output Fields

Table 130 on page 1872 describes the output fields for the **show multicast sessions** command. Output fields are listed in the approximate order in which they appear.

Table 130: show multicast sessions Output Fields

Field Name	Field Description
session-name	Name of the known announced multicast sessions.

 Copyright © 2018, Juniper Networks, Inc.

Sample Output

show multicast sessions

```
user@host> show multicast sessions
1-Department of Biological Sciences, LSU
  ...
  Monterey Bay - DockCam
  Monterey Bay - JettyCam
  Monterey Bay - StandCam
  Monterey DockCam
  Monterey DockCam / ROV cam
  ...
  NASA TV (MPEG-1)
  ...
  UO Broadcast - NASA Videos - 25 Years of Progress
  UO Broadcast - NASA Videos - Journey through the Solar System
  UO Broadcast - NASA Videos - Life in the Universe
  UO Broadcast - NASA Videos - Nasa and the Airplane
  UO Broadcasts OPB's Oregon Story
  UO DOD News Clips
  UO Medical Management of Biological Casualties (1)
  UO Medical Management of Biological Casualties (2)
  UO Medical Management of Biological Casualties (3)
...
  376 active sessions.
```

show multicast sessions regular-expression detail

```
user@host> show multicast sessions "NASA TV" detail
SDP Version: 0  Originated by: -@10.223.83.33
 Session: NASA TV (MPEG-1)
 Description: NASA television in MPEG-1 format, provided by Private University.
Please contact the UO if you have problems with this feed.
 Email: Your Name Here <multicast@lists.private.edu>
 Phone: Your Name Here <888/555-1212>
Bandwidth: AS:1000
Start time: permanent
Stop time: none
Attribute: type:broadcast
Attribute: tool:IP/TV Content Manager 3.4.14
Attribute: live:capture:1
Attribute: x-iptv-capture:mp1s
Media: video 54302 RTP/AVP 32 31 96 97
Connection Data: 233.252.0.45 ttl 127
Attribute: quality:8
Attribute: framerate:30
Attribute: rtpmap:96 WBIH/90000
Attribute: rtpmap:97 MP4V-ES/90000
Attribute: x-iptv-svr:video 10.223.91.191 live
Attribute: fmtp:32 type=mpeg1
Media: audio 28848 RTP/AVP 14 0 96 3 5 97 98 99 100 101 102 10 11 103 104 105 106
 Connection Data: 224.2.145.37 ttl 127
Attribute: rtpmap:96 X-WAVE/8000
Attribute: rtpmap:97 L8/8000/2
Attribute: rtpmap:98 L8/8000
Attribute: rtpmap:99 L8/22050/2
Attribute: rtpmap:100 L8/22050
Attribute: rtpmap:101 L8/11025/2
Attribute: rtpmap:102 L8/11025
```

Copyright © 2018, Juniper Networks, Inc.

```
Attribute: rtpmap:103 L16/22050/2
Attribute: rtpmap:104 L16/22050

  1 matching sessions.
```


Copyright © 2018, Juniper Networks, Inc.

show multicast snooping route

Syntax

show multicast snooping route
<*regexp*>
<active>
<all>
<bridge-domain *bridge-domain-name*>
<brief >
<control>
<data>
<detail >
<extensive>
<group *group*>
<inactive>
<inet>
<inet6>
<instance *instance-name*>
<logical-system *logical-system-name*>
<mesh-group *mesh-group-name*>
<qualified-vlan *vlan-id*>
<source-prefix *source-prefix*>
<vlan *vlan-id*>

Release Information

Command introduced in Junos OS Release 8.5.
Support for **control**, **data**, **qualified-vlan** and **vlan** options introduced in Junos OS Release 13.3 for EX Series switches.

Description

Display the entries in the IP multicast snooping forwarding table. You can display some of this information with the **show route table inet.1** command.

Options

none—Display standard information about all entries in the multicast snooping table for all virtual switches and all bridge domains.

active | all | inactive —(Optional) Display all active entries, all entries, or all inactive entries, respectively, in the multicast snooping table.

bridge-domain *bridge-domain*—(Optional) Display the entries for a particular bridge domain.

brief | detail | extensive—(Optional) Display the specified level of output.

control—(Optional) Display control route entries.

data—(Optional) Display data route entries.

group *group*—(Optional) Display the entries for a particular group.

inet—(Optional) Display IPv4 information.

inet6—(Optional) Display IPv6 information.

instance *instance-name*—(Optional) Display the entries for a multicast instance.

Copyright © 2018, Juniper Networks, Inc.

logical-system ***logical-system-name***—(Optional) Display information about a particular logical system, or type 'all'.

mesh-group ***mesh-group-name***—(Optional) Display the entries for a particular mesh group.

qualified-vlan ***vlan-id***—(Optional) Display the entries for a particular qualified VLAN.

regexp—(Optional) Display information about the multicast forwarding table entries that match a UNIX-style regular expression.

source-prefix ***source-prefix***—(Optional) Display the entries for a particular source prefix.

vlan ***vlan-id***—(Optional) Display the entries for a particular VLAN.

Required Privilege Level view

List of Sample Output

show multicast snooping route bridge-domain on page 1877
show multicast snooping route instance vs on page 1877
show multicast snooping route extensive on page 1877

Output Fields Table 131 on page 1876 describes the output fields for the **show multicast snooping route** command. Output fields are listed in the approximate order in which they appear.

Table 131: show multicast snooping route Output Fields

Field Name	Field Description	Level of Output
Nexthop Bulking	Displays whether next-hop bulk updating is **ON** or **OFF** (only for routing-instance type of **virtual switch** or **vpls**).	All levels
Family	IPv4 address family (**INET**) or IPv6 address family (**INET6**).	All levels
Group	Group address.	All levels
Source	Prefix and length of the source as it is in the multicast forwarding table. For (*,G) entries, this field is set to "*".	All levels
Routing-instance	Name of the routing instance to which this routing information applies. (Displayed when multicast is configured within a routing instance.)	All levels
Learning Domain	Name of the learning domain to which this routing information applies.	**detail extensive**
Statistics	Rate at which packets are being forwarded for this source and group entry (in Kbps and pps), and number of packets that have been forwarded to this prefix.	**detail extensive**
Next-hop ID	Next-hop identifier of the prefix. The identifier is returned by the router's Packet Forwarding Engine and is also displayed in the output of the **show multicast nexthops** command.	**detail extensive**
Route state	Whether the group is **Active** or **Inactive**.	**extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 131: show multicast snooping route Output Fields (continued)

Field Name	Field Description	Level of Output
Forwarding state	Whether the prefix is **Pruned** or **Forwarding**.	**extensive**
Cache lifetime/timeout	Number of seconds until the prefix is removed from the multicast forwarding table. A value of **never** indicates a permanent forwarding entry.	**extensive**

Sample Output

show multicast snooping route bridge-domain

```
user@host> show multicast snooping route bridge-domain br-dom-1 extensive
Family: INET

Group: 232.1.1.1
    Source: 192.168.3.100/32
    Downstream interface list:
        ge-0/1/0.200
    Statistics: 0 kBps, 0 pps, 1 packets
    Next-hop ID: 1048577
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: 240 seconds
```

show multicast snooping route instance vs

```
user@host> show multicast snooping route instance vs
Nexthop Bulking: ON

Family: INET

Group: 224.0.0.0
    Bridge-domain: vsid500

Group: 225.1.0.1
    Bridge-domain: vsid500
    Downstream interface list: vsid500
       ge-0/3/8.500 ge-1/1/9.500 ge1/2/5.500
```

show multicast snooping route extensive

```
user@host> show multicast snooping route extensive inet6 group ff03::1
Nexthop Bulking: OFF

Family: INET6
Group: ff03::1/128
    Source: ::
    Bridge-domain: BD-1
    Mesh-group: __all_ces__
    Downstream interface list:
        ae0.1 -(562) 1048576
    Statistics: 2697 kBps, 3875 pps, 758819039 packets
    Next-hop ID: 1048605
    Route state: Active
    Forwarding state: Forwarding
```

Copyright © 2018, Juniper Networks, Inc.

```
Group: ff03::1/128
    Source: 6666::2/128
    Bridge-domain: BD-1
    Mesh-group: __all_ces__
    Downstream interface list:
        ae0.1 -(562) 1048576
    Statistics: 0 kBps, 0 pps, 0 packets
    Next-hop ID: 1048605
    Route state: Active
    Forwarding state: Forwarding
```

 Copyright © 2018, Juniper Networks, Inc.

show multicast statistics

Syntax

show multicast statistics
<inet | inet6>
<instance *instance-name*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced before Junos OS Release 7.4.
interface option introduced in Junos OS Release 16.1 for the MX Series.

Description

Display IP multicast statistics.

Options

none—Display multicast statistics for all supported address families for all routing instances.

inet | inet6—(Optional) Display multicast statistics for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display statistics for a specific routing instance.

interface *interface-name*—(Optional) Display statistics for a specific interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information

The input and output interface multicast statistics are consistent, but not timely. They are constructed from the forwarding statistics, which are gathered at 30-second intervals. Therefore, the output from this command always lags the true count by up to 30 seconds.

Required Privilege Level

view

Related Documentation

- clear multicast statistics on page 1748

List of Sample Output

show multicast statistics on page 1882
show multicast statistics interface on page 1882

Output Fields

Table 132 on page 1879 describes the output fields for the **show multicast statistics** command. Output fields are listed in the approximate order in which they appear.

Table 132: show multicast statistics Output Fields

Field Name	Field Description
Family	Protocol family for which multicast statistics are displayed: **INET** or **INET6**.

Copyright © 2018, Juniper Networks, Inc.

Table 132: show multicast statistics Output Fields (continued)

Field Name	Field Description
Interface	Name of the interface for which statistics are being reported.
Routing Protocol	Primary multicast protocol on the interface: **PIM**, **DVMRP for INET**, or **PIM for INET6**.
Mismatch	Number of multicast packets that did not arrive on the correct upstream interface.
Kernel Resolve	Number of resolve requests processed by the primary multicast protocol on the interface.
Resolve No Route	Number of resolve requests that were ignored because there was no route to the source.
In Kbytes	Total accumulated incoming packets (in KB) since the last time the **clear multicast statistics** command was issued.
Out Kbytes	Total accumulated outgoing packets (in KB) since the last time the **clear multicast statistics** command was issued.
Mismatch error	Number of mismatches that were ignored because of internal errors.
Mismatch No Route	Number of mismatches that were ignored because there was no route to the source.
Routing Notify	Number of times that the multicast routing system has been notified of a new multicast source by a multicast routing protocol .
Resolve Error	Number of resolve requests that were ignored because of internal errors.
In Packets	Total number of incoming packets since the last time the **clear multicast statistics** command was issued.
Out Packets	Total number of outgoing packets since the last time the **clear multicast statistics** command was issued.
Resolve requests on interfaces not enabled for multicast *n*	Number of resolve requests on interfaces that are not enabled for multicast that have accumulated since the **clear multicast statistics** command was last issued.
Resolve requests with no route to source *n*	Number of resolve requests with no route to the source that have accumulated since the **clear multicast statistics** command was last issued.
Routing notifications on interfaces not enabled for multicast *n*	Number of routing notifications on interfaces not enabled for multicast that have accumulated since the **clear multicast statistics** command was last issued.
Routing notifications with no route to source *n*	Number of routing notifications with no route to the source that have accumulated since the **clear multicast statistics** command was last issued.
Interface Mismatches on interfaces not enabled for multicast *n*	Number of interface mismatches on interfaces not enabled for multicast that have accumulated since the **clear multicast statistics** command was last issued.

Copyright © 2018, Juniper Networks, Inc.

Table 132: show multicast statistics Output Fields (continued)

Field Name	Field Description
Group Membership on interfaces not enabled for multicast *n*	Number of group memberships on interfaces not enabled for multicast that have accumulated since the **clear multicast statistics** command was last issued.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show multicast statistics

```
user@host> show multicast statistics
    Address family: INET
    Interface: fe-0/0/0
        Routing Protocol:          PIM   Mismatch error:               0
        Mismatch:                    0   Mismatch No Route:            0
        Kernel Resolve:             10   Routing Notify:               0
        Resolve No Route:            0   Resolve Error:                0
        In Kbytes:                4641   In Packets:               50454
        Out Kbytes:                  0   Out Packets:                  0
    Interface: so-0/1/1.0
        Routing Protocol:          PIM   Mismatch error:               0
        Mismatch:                    0   Mismatch No Route:            0
        Kernel Resolve:              0   Routing Notify:               0
        Resolve No Route:            0   Resolve Error:                0
        In Kbytes:                   0   In Packets:                   0
        Out Kbytes:               4641   Out Packets:              50454

    Resolve requests on interfaces not enabled for multicast 0
    Resolve requests with no route to source 0
    Routing notifications on interfaces not enabled for multicast 0
    Routing notifications with no route to source 0
    Interface Mismatches on interfaces not enabled for multicast 0
    Group Membership on interfaces not enabled for multicast 25

    Address family: INET6
    Interface: fe-0/0/0.0
        Routing Protocol:          PIM   Mismatch error:               0
        Mismatch:                    0   Mismatch No Route:            0
        Kernel Resolve:              0   Routing Notify:               0
        Resolve No Route:            0   Resolve Error:                0
        In Kbytes:                   0   In Packets:                   0
        Out Kbytes:                  0   Out Packets:                  0
    Interface: so-0/1/1.0
        Routing Protocol:          PIM   Mismatch error:               0
        Mismatch:                    0   Mismatch No Route:            0
        Kernel Resolve:              0   Routing Notify:               0
        Resolve No Route:            0   Resolve Error:                0
        In Kbytes:                   0   In Packets:                   0
        Out Kbytes:                  0   Out Packets:                  0

    Resolve requests on interfaces not enabled for multicast 0
    Resolve requests with no route to source 0
    Routing notifications on interfaces not enabled for multicast 0
    Routing notifications with no route to source 0
    Interface Mismatches on interfaces not enabled for multicast 0
    Group Membership on interfaces not enabled for multicast 0
```

show multicast statistics interface

```
user@host> show multicast statistics interface vt-3/0/10.2097152
Instance: master Family: INET
Interface: vt-3/0/10.2097152
    Routing protocol:          PIM   Mismatch error:               0
    Mismatch:                    0   Mismatch no route:            0
    Kernel resolve:              0   Routing notify:               0
    Resolve no route:            0   Resolve error:                0
```

 Copyright © 2018, Juniper Networks, Inc.

```
Resolve filtered:             0   Notify filtered:              0
In kbytes:                    0   In packets:                   0
Out kbytes:                   0   Out packets:                  0
```

show multicast usage

List of Syntax

Syntax on page 1884
Syntax (EX Series Switch and the QFX Series) on page 1884

Syntax

show multicast usage
<brief | detail>
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)

show multicast usage
<brief | detail>
<inet | inet6>
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display usage information about the 10 most active Distance Vector Multicast Routing Protocol (DVMRP) or Protocol Independent Multicast (PIM) groups.

Options

none—Display multicast usage information for all supported address families for all routing instances.

brief | detail—(Optional) Display the specified level of output.

inet | inet6—(Optional) Display usage information for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display information about the most active DVMRP or PIM groups for a specific multicast instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show multicast usage on page 1885
show multicast usage brief on page 1885
show multicast usage instance on page 1885
show multicast usage detail on page 1886

Output Fields

Table 133 on page 1885 describes the output fields for the **show multicast usage** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2018, Juniper Networks, Inc.

Table 133: show multicast usage Output Fields

Field Name	Field Description
Instance	Name of the routing instance. (Displayed when multicast is configured within a routing instance.)
Group	Group address.
Sources	Number of sources.
Packets	Number of packets that have been forwarded to this prefix. If one or more of the packets forwarded statistic queries fails or times out, the packets field displays **unavailable**.
Bytes	Number of bytes that have been forwarded to this prefix. If one or more of the packets forwarded statistic queries fails or times out, the bytes field displays **unavailable**.
Prefix	IP address.
/len	Prefix length.
Groups	Number of multicast groups.

Sample Output

show multicast usage

```
user@host> show multicast usage
Group              Sources  Packets               Bytes
233.252.0.0          1        52847                 4439148
233.252.0.1          2        13450                 1125530


Prefix             /len  Groups  Packets               Bytes
10.255.14.144      /32   2       66254                 5561304
10.255.70.15       /32   1       43                    3374...
```

show multicast usage brief

The output for the **show multicast usage brief** command is identical to that for the **show multicast usage** command. For sample output, see show multicast usage on page 1885.

show multicast usage instance

```
user@host> show multicast usage instance VPN-A
Group              Sources  Packets               Bytes
233.252.0.254      1        5538                  509496
233.252.0.39       1        13                    624
233.252.0.40       1        13                    624

Prefix             /len  Groups  Packets               Bytes
192.168.195.34     /32   1       5538                  509496
```

```
10.255.14.30    /32  1      13                  624
10.255.245.91   /32  1      13                  624
...
```

show multicast usage detail

```
user@host> show multicast usage detail
Group           Sources Packets              Bytes
233.252.0.0       1       53159                4465356
    Source: 10.255.14.144   /32 Packets: 53159 Bytes: 4465356
233.252.0.1       2       13450                1125530
    Source: 10.255.14.144   /32 Packets: 13407 Bytes: 1122156
    Source: 10.255.70.15    /32 Packets: 43 Bytes: 3374


Prefix          /len Groups Packets              Bytes
10.255.14.144   /32  2      66566                5587512
    Group: 233.252.0.0       Packets: 53159 Bytes: 4465356
    Group: 233.252.0.1       Packets: 13407 Bytes: 1122156
10.255.70.15    /32  1      43                   3374
    Group: 233.252.0.1       Packets: 43 Bytes: 3374
```

 Copyright © 2018, Juniper Networks, Inc.

show pim bootstrap

List of Syntax

Syntax on page 1887
Syntax (EX Series Switch and the QFX Series) on page 1887

Syntax

show pim bootstrap
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)

show pim bootstrap
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
instance option introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

For sparse mode only, display information about Protocol Independent Multicast (PIM) bootstrap routers.

Options

none—Display PIM bootstrap router information for all routing instances.

instance *instance-name*—(Optional) Display information about bootstrap routers for a specific PIM-enabled routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show pim bootstrap on page 1888
show pim bootstrap instance on page 1888

Output Fields

Table 134 on page 1887 describes the output fields for the **show pim bootstrap** command. Output fields are listed in the approximate order in which they appear.

Table 134: show pim bootstrap Output Fields

Field Name	Field Description
Instance	Name of the routing instance.
BSR	Bootstrap router.
Pri	Priority of the routing device as elected to be the bootstrap router.

Copyright © 2018, Juniper Networks, Inc.

Table 134: show pim bootstrap Output Fields (continued)

Field Name	Field Description
Local address	Local routing device address.
Pri	Local routing device address priority to be elected as the bootstrap router.
State	Local routing device election state: **Candidate**, **Elected**, or **Ineligible**.
Timeout	How long until the local routing device declares the bootstrap router to be unreachable, in seconds.

Sample Output

show pim bootstrap

```
user@host> show pim bootstrap
Instance: PIM.master

BSR                     Pri Local address           Pri State      Timeout
None                      0 10.255.71.46              0 InEligible       0
2001:db8:1:1:1:0:aff:785c 34 2001:db8:1:1:1:0:aff:7c12   0 InEligible       0
```

show pim bootstrap instance

```
user@host> show pim bootstrap instance VPN-A
Instance: PIM.VPN-A

BSR              Pri Local address   Pri State      Timeout
None               0 192.168.196.105   0 InEligible       0
```


Copyright © 2018, Juniper Networks, Inc.

show pim interfaces

List of Syntax

Syntax on page 1889
Syntax (EX Series Switch and the QFX Series) on page 1889

Syntax

show pim interfaces
<inet | inet6>
<instance (*instance-name* | all)>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)

show pim interfaces
<inet | inet6>
<instance (*instance-name* | all)>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Commmand introduced in Junos OS Release 11.3 for the QFX Series.
Support for bidirectional PIM added in Junos OS Release 12.1.
Support for the **instance all** option added in Junos OS Release 12.1.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display information about the interfaces on which Protocol Independent Multicast (PIM) is configured.

Options

none—Display interface information for all family addresses for the main instance.

inet | inet6—(Optional) Display interface information for IPv4 or IPv6 family addresses, respectively.

instance (*instance-name* | all)—(Optional) Display information about interfaces for a specific PIM-enabled routing instance or for all routing instances.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show pim interfaces on page 1891

Output Fields

Table 135 on page 1889 describes the output fields for the **show pim interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 135: show pim interfaces Output Fields

Field Name	Field Description
Instance	Name of the routing instance.

Copyright © 2018, Juniper Networks, Inc.

Table 135: show pim interfaces Output Fields (continued)

Field Name	Field Description
Name	Interface name.
State	State of the interface. The state also is displayed in the **show interfaces** command.
Mode	PIM mode running on the interface: • **B**—In bidirectional mode, multicast groups are carried across the network over bidirectional shared trees. This type of tree minimizes PIM routing state, which is especially important in networks with numerous and dispersed senders and receivers. • **S**—In sparse mode, routing devices must join and leave multicast groups explicitly. Upstream routing devices do not forward multicast traffic to this routing device unless this device has sent an explicit request (using a join message) to receive multicast traffic. • **Dense**—Unlike sparse mode, where data is forwarded only to routing devices sending an explicit request, dense mode implements a flood-and-prune mechanism, similar to DVMRP (the first multicast protocol used to support the multicast backbone). (Not supported on QFX Series.) • **Sparse-Dense**—Sparse-dense mode allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as **dense** is not mapped to a rendezvous point (RP). Instead, data packets destined for that group are forwarded using PIM-Dense Mode (PIM-DM) rules. A group specified as **sparse** is mapped to an RP, and data packets are forwarded using PIM-Sparse Mode (PIM-SM) rules. When sparse-dense mode is configured, the output includes both **S** and **D**. When bidirectional-sparse mode is configured, the output includes **S** and **B**. When bidirectional-sparse-dense mode is configured, the output includes **B**, **S**, and **D**.
IP	Version number of the address family on the interface: **4** (IPv4) or **6** (IPv6).
V	PIM version running on the interface: **1** or **2**.
State	State of PIM on the interface: • **Active**—Bidirectional mode is enabled on the interface and on all PIM neighbors. • **DR**—Designated router. • **NotCap**—Bidirectional mode is not enabled on the interface. This can happen when bidirectional PIM is not configured locally, when one of the neighbors is not configured for bidirectional PIM, or when one of the neighbors has not implemented the bidirectional PIM protocol. • **NotDR**—Not the designated router. • **P2P**—Point to point.
NbrCnt	Number of neighbors that have been seen on the interface.
JoinCnt(sg)	Number of (s,g) join messages that have been seen on the interface.
JointCnt(*g)	Number of (*g) join messages that have been seen on the interface.
DR address	Address of the designated router.

 Copyright © 2018, Juniper Networks, Inc.

Sample Output

show pim interfaces

```
user@host> show pim interfaces
Stat = Status, V = Version, NbrCnt = Neighbor Count,
S = Sparse, D = Dense, B = Bidirectional,
DR = Designated Router, P2P = Point-to-point link,
Active = Bidirectional is active, NotCap = Not Bidirectional Capable

Name               Stat Mode IP V State        NbrCnt JoinCnt(sg/*g) DR address
ge-0/3/0.0         Up   S     4 2 NotDR,NotCap      1 0/0            40.0.0.3
ge-0/3/3.50        Up   S     4 2 DR,NotCap         1 9901/100       50.0.0.2
ge-0/3/3.51        Up   S     4 2 DR,NotCap         1 0/0            51.0.0.2
pe-1/2/0.32769     Up   S     4 2 P2P,NotCap        0 0/0
```

show pim join

List of Syntax

Syntax on page 1892
Syntax (EX Series Switch and the QFX Series) on page 1892

Syntax

show pim join
<brief | detail | extensive | summary>
<bidirectional | dense | sparse>
<downstream-count>
<exact>
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<*range*>
<rp *ip-address/prefix* | source *ip-address/prefix*>
<sg | star-g>

Syntax (EX Series Switch and the QFX Series)

show pim join
<brief | detail | extensive | summary>
<dense | sparse>
<exact>
<inet | inet6>
<instance *instance-name*>
<*range*>
<rp *ip-address/prefix* | source *ip-address/prefix*>
<sg | star-g>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
summary option introduced in Junos OS Release 9.6.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Support for bidirectional PIM added in Junos OS Release 12.1.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Multiple new filter options introduced in Junos OS Release 13.2.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
downstream-count option introduced in Junos OS Release 16.1.
Support for PIM NSR support for VXLAN added in Junos OS Release 16.2
Support for RFC 5496 (via **rpf-vector**) added in Junos OS Release 17.3R1.

Description

Display information about Protocol Independent Multicast (PIM) groups for all PIM modes.

For bidirectional PIM, display information about PIM group ranges (*,G-range) for each active bidirectional RP group range, in addition to each of the joined (*,G) routes.

Options

none—Display the standard information about PIM groups for all supported family addresses for all routing instances.

brief | detail | extensive | summary—(Optional) Display the specified level of output.

 Copyright © 2018, Juniper Networks, Inc.

bidirectional | dense | sparse—(Optional) Display information about PIM bidirectional mode, dense mode, or sparse and source-specific multicast (SSM) mode entries.

downstream-count—(Optional) Display the downstream count instead of a list.

exact—(Optional) Display information about only the group that exactly matches the specified group address.

inet | inet6—(Optional) Display PIM group information for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display information about groups for the specified PIM-enabled routing instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

range—(Optional) Address range of the group, specified as ***prefix/prefix-length***.

rp *ip-address/prefix* | source *ip-address/prefix*—(Optional) Display information about the PIM entries with a specified rendezvous point (RP) address and prefix or with a specified source address and prefix. You can omit the prefix.

sg | star-g—(Optional) Display information about PIM (S,G) or (*,G) entries.

Required Privilege Level view

Related Documentation

- clear pim join on page 1750
- *Example: Configuring Multicast-Only Fast Reroute in a PIM Domain*
- *Example: Configuring Bidirectional PIM*
- *Example: Configuring PIM State Limits*

List of Sample Output

Copyright © 2018, Juniper Networks, Inc.

show pim join extensive (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 1904
show pim join extensive (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 1905

Output Fields Table 136 on page 1894 describes the output fields for the **show pim join** command. Output fields are listed in the approximate order in which they appear.

Table 136: show pim join Output Fields

Field Name	Field Description	Level of Output
Instance	Name of the routing instance.	**brief detail extensive summary** none
Family	Name of the address family: **inet** (IPv4) or **inet6** (IPv6).	**brief detail extensive summary** none
Route type	Type of multicast route: (S,G) or (*,G).	**summary**
Route count	Number of (S,G) routes and number of (*,G) routes.	**summary**
R	Rendezvous Point Tree.	**brief detail extensive** none
S	Sparse.	**brief detail extensive** none
W	Wildcard.	**brief detail extensive** none
Group	Group address.	**brief detail extensive** none
Bidirectional group prefix length	For bidirectional PIM, length of the IP prefix for RP group ranges.	All levels
Source	Multicast source: • ***** (wildcard value) • ***ipv4-address*** • ***ipv6-address***	**brief detail extensive** none
RP	Rendezvous point for the PIM group.	**brief detail extensive** none
Flags	PIM flags: • **bidirectional**—Bidirectional mode entry. • **dense**—Dense mode entry. • **rptree**—Entry is on the rendezvous point tree. • **sparse**—Sparse mode entry. • **spt**—Entry is on the shortest-path tree for the source. • **wildcard**—Entry is on the shared tree.	**brief detail extensive** none

 Copyright © 2018, Juniper Networks, Inc.

Table 136: show pim join Output Fields (continued)

Field Name	Field Description	Level of Output
Upstream interface	RPF interface toward the source address for the source-specific state (S,G) or toward the rendezvous point (RP) address for the non-source-specific state (*,G). For bidirectional PIM, **RP Link** means that the interface is directly connected to a subnet that contains a phantom RP address. A pseudo multipoint LDP (M-LDP) interface appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.	**brief detail extensive** none
Upstream neighbor	Information about the upstream neighbor: **Direct**, **Local**, **Unknown**, or a specific IP address. For bidirectional PIM, **Direct** means that the interface is directly connected to a subnet that contains a phantom RP address. The multipoint LDP (M-LDP) root appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.	**extensive**
Upstream rpf-vector	Information about the upstream Reverse Path Forwarding (RPF) vector; appears in conjunction with the **rpf-vector** command.	**extensive**
Active upstream interface	When multicast-only fast reroute (MoFRR) is configured in a PIM domain, the upstream interface for the active path. A PIM router propagates join messages on two upstream RPF interfaces to receive multicast traffic on both links for the same join request. Preference is given to two paths that do not converge to the same immediate upstream router. PIM installs appropriate multicast routes with upstream neighbors as RPF next hops with two (primary and backup) interfaces.	**extensive**
Active upstream neighbor	On the MoFRR primary path, the IP address of the neighbor that is directly connected to the active upstream interface.	**extensive**
MoFRR Backup upstream interface	The MoFRR upstream interface that is used when the primary path fails. When the primary path fails, the backup path is upgraded to primary, and traffic is forwarded accordingly. If there are alternate paths available, a new backup path is calculated and the appropriate multicast route is updated or installed.	**extensive**
MoFRR Backup upstream neighbor	IP address of the MoFRR upstream neighbor.	extensive

Copyright © 2018, Juniper Networks, Inc.

Table 136: show pim join Output Fields (continued)

Field Name	Field Description	Level of Output
Upstream state	Information about the upstream interface: • **Join to RP**—Sending a join to the rendezvous point. • **Join to Source**—Sending a join to the source. • **Local RP**—Sending neither join messages nor prune messages toward the RP, because this routing device is the rendezvous point. • **Local Source**—Sending neither join messages nor prune messages toward the source, because the source is locally attached to this routing device. • **No Prune to RP**—Automatically sent to RP when SPT and RPT are on the same path. • **Prune to RP**—Sending a prune to the rendezvous point. • **Prune to Source**—Sending a prune to the source. NOTE: RP group range entries have **None** in the **Upstream state** field because RP group ranges do not trigger actual PIM join messages between routing devices.	**extensive**
Downstream neighbors	Information about downstream interfaces: • **Interface**—Interface name for the downstream neighbor. A pseudo PIM-SM interface appears for all IGMP-only interfaces. A pseudo multipoint LDP (Pseudo-MLDP) interface appears on ingress root nodes in M-LDP point-to-multipoint LSPs with inband signaling. • **Interface address**—Address of the downstream neighbor. • **State**—Information about the downstream neighbor: **join** or **prune**. • **Flags**—PIM join flags: **R (RPtree)**, **S (Sparse)**, **W (Wildcard)**, or **zero**. • **Uptime**—Time since the downstream interface joined the group. • **Time since last Join**—Time since the last join message was received from the downstream interface. • **Time since last Prune**—Time since the last prune message was received from the downstream interface. • **rpf-vector**—IP address of the RPF vector TLV .	**extensive**
Number of downstream interfaces	Total number of outgoing interfaces for each (S,G) entry.	**extensive**
Assert Timeout	Length of time between assert cycles on the downstream interface. Not displayed if the assert timer is null.	**extensive**
Keepalive timeout	Time remaining until the downstream join state is updated (in seconds). If the downstream join state is not updated before this keepalive timer reaches zero, the entry is deleted. If there is a directly connected host, **Keepalive timeout** is **Infinity**.	**extensive**
Uptime	Time since the creation of (S,G) or (*,G) state. The uptime is not refreshed every time a PIM join message is received for an existing (S,G) or (*,G) state.	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 136: show pim join Output Fields (continued)

Field Name	Field Description	Level of Output
Bidirectional accepting interfaces	Interfaces on the routing device that forward bidirectional PIM traffic. The reasons for forwarding bidirectional PIM traffic are that the interface is the winner of the designated forwarder election (**DF Winner**), or the interface is the reverse path forwarding (RPF) interface toward the RP (**RPF**).	**extensive**

Sample Output

show pim join summary

```
user@host> show pim join summary
Instance: PIM.master Family: INET

Route type              Route count
(s,g)                   2
(*,g)                   1

Instance: PIM.master Family: INET6
```

show pim join (PIM Sparse Mode)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 233.252.0.1
    Source: *
    RP: 10.255.14.144
    Flags: sparse,rptree,wildcard
    Upstream interface: Local

Group: 233.252.0.1
    Source: 10.255.14.144
    Flags: sparse,spt
    Upstream interface: Local

Group: 233.252.0.1
    Source: 10.255.70.15
    Flags: sparse,spt
    Upstream interface: so-1/0/0.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join (Bidirectional PIM)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 233.252.0.1
    Bidirectional group prefix length: 24
    Source: *
```

Copyright © 2018, Juniper Networks, Inc.

```
        RP: 10.10.13.2
        Flags: bidirectional,rptree,wildcard
        Upstream interface: ge-0/0/1.0

    Group: 233.252.0.2
        Bidirectional group prefix length: 24
        Source: *
        RP: 10.10.1.3
        Flags: bidirectional,rptree,wildcard
        Upstream interface: ge-0/0/1.0 (RP Link)

    Group: 233.252.0.3
        Bidirectional group prefix length: 24
        Source: *
        RP: 10.10.13.2
        Flags: bidirectional,rptree,wildcard
        Upstream interface: ge-0/0/1.0

    Group: 233.252.0.4
        Bidirectional group prefix length: 24
        Source: *
        RP: 10.10.1.3
        Flags: bidirectional,rptree,wildcard
        Upstream interface: ge-0/0/1.0 (RP Link)

    Instance: PIM.master Family: INET6
    R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join inet6

```
user@host> show pim join inet6
Instance: PIM.master Family: INET6
 R = Rendezvous Point Tree, S = Sparse, W = Wildcard

 Group: 2001:db8::e000:101
     Source: *
     RP: ::46.0.0.13
     Flags: sparse,rptree,wildcard
     Upstream interface: Local

 Group: 2001:db8::e000:101
     Source: ::1.1.1.1
     Flags: sparse
     Upstream interface: unknown (no neighbor)

 Group: 2001:db8::e800:101
     Source: ::1.1.1.1
     Flags: sparse
     Upstream interface: unknown (no neighbor)

 Group: 2001:db8::e800:101
     Source: ::1.1.1.2
     Flags: sparse
     Upstream interface: unknown (no neighbor)
```

show pim join inet6 star-g

```
user@host> show pim join inet6 star-g
Instance: PIM.master Family: INET6
 R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

 Copyright © 2018, Juniper Networks, Inc.

```
Group: 2001:db8::e000:101
    Source: *
    RP: ::46.0.0.13
    Flags: sparse,rptree,wildcard
    Upstream interface: Local
```

show pim join instance <instance-name>

```
user@host> show pim join instance VPN-A
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 233.252.0.2
    Source: *
    RP: 10.10.47.100
    Flags: sparse,rptree,wildcard
    Upstream interface: Local

Group: 233.252.0.2
    Source: 192.168.195.74
    Flags: sparse,spt
    Upstream interface: at-0/3/1.0

Group: 233.252.0.2
    Source: 192.168.195.169
    Flags: sparse
    Upstream interface: so-1/0/1.0

Instance: PIM.VPN-A Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join instance <instance-name> downstream-count

```
user@host> show pim join instance VPN-A downstream-count
Instance: PIM.SML_VRF_4 Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 233.252.0.1
    Source: *
    RP: 10.11.11.6
    Flags: sparse,rptree,wildcard
    Upstream interface: mt-1/2/10.32813
    Number of downstream interfaces: 4

Group: 233.252.0.1
    Source: 10.1.1.1
    Flags: sparse,spt
    Upstream interface: ge-0/0/3.5
    Number of downstream interfaces: 5
```

show pim join instance <instance-name> downstream-count extensive

```
user@host> show pim join instance VPN-A downstream-count extensive
Instance: PIM.SML_VRF_4 Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 233.252.0.1
    Source: *
    RP: 10.11.11.6
```

Copyright © 2018, Juniper Networks, Inc.

```
    Flags: sparse,rptree,wildcard
    Upstream interface: mt-1/2/10.32813
    Upstream neighbor: 10.2.2.7 (assert winner)
    Upstream state: Join to RP
    Uptime: 02:51:41
    Number of downstream interfaces: 4
    Number of downstream neighbors: 4

Group: 233.252.0.1
    Source: 10.1.1.1
    Flags: sparse,spt
    Upstream interface: ge-0/0/3.5
    Upstream neighbor: 10.1.1.17
    Upstream state: Join to Source, Prune to RP
    Keepalive timeout: 0
    Uptime: 02:51:42
    Number of downstream interfaces: 5
    Number of downstream neighbors: 7
```

show pim join detail

```
user@host> show pim join detail
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 233.252.0.1
    Source: *
    RP: 10.255.14.144
    Flags: sparse,rptree,wildcard
    Upstream interface: Local

Group: 233.252.0.1
    Source: 10.255.14.144
    Flags: sparse,spt
    Upstream interface: Local

Group: 233.252.0.1
    Source: 10.255.70.15
    Flags: sparse,spt
    Upstream interface: so-1/0/0.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join extensive (PIM Resolve TLV for Multicast in Seamless MPLS)

```
user@host> show pim join extensive
Group: 228.26.1.5
    Source: 60.0.0.101
    Flags: sparse,spt
    Upstream interface: ge-5/0/0.1
    Upstream neighbor: 10.100.1.13
    Upstream state: Join to Source
Upstream rpf-vector: 10.100.20.1
    Keepalive timeout: 178
    Uptime: 17:44:38
    Downstream neighbors:
        Interface: xe-2/0/3.1
            203.21.2.190 State: Join Flags: S Timeout: 156
            Uptime: 17:44:38 Time since last Join: 00:00:54
```

 Copyright © 2018, Juniper Networks, Inc.

```
                **rpf-vector: 10.100.20.1**
            Interface: xe-2/0/2.1
                203.21.1.190 State: Join Flags: S Timeout: 156
                Uptime: 17:44:38 Time since last Join: 00:00:54
                **rpf-vector: 10.100.20.2**
        Number of downstream interfaces: 2
        Number of downstream neighbors: 2
```

show pim join extensive (PIM Sparse Mode)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 233.252.0.1
    Source: *
    RP: 10.255.14.144
    Flags: sparse,rptree,wildcard
    Upstream interface: Local
    Upstream neighbor: Local
    Upstream state: Local RP
    Uptime: 00:03:49
    Downstream neighbors:
        Interface: so-1/0/0.0
            10.111.10.2 State: Join Flags: SRW Timeout: 174
            Uptime: 00:03:49 Time since last Join: 00:01:49
        Interface: mt-1/1/0.32768
            10.10.47.100 State: Join Flags: SRW  Timeout: Infinity
            Uptime: 00:03:49 Time since last Join: 00:01:49
    Number of downstream interfaces: 2

Group: 233.252.0.1
    Source: 10.255.14.144
    Flags: sparse,spt
    Upstream interface: Local
    Upstream neighbor: Local
    Upstream state: Local Source, Local RP
    Keepalive timeout: 344
    Uptime: 00:03:49
    Downstream neighbors:
        Interface: so-1/0/0.0
            10.111.10.2 State: Join Flags: S Timeout: 174
            Uptime: 00:03:49 Time since last Prune: 00:01:49
        Interface: mt-1/1/0.32768
            10.10.47.100 State: Join Flags: S   Timeout: Infinity
            Uptime: 00:03:49 Time since last Prune: 00:01:49
    Number of downstream interfaces: 2

Group: 233.252.0.1
    Source: 10.255.70.15
    Flags: sparse,spt
    Upstream interface: so-1/0/0.0
    Upstream neighbor: 10.111.10.2
    Upstream state: Local RP, Join to Source
    Keepalive timeout: 344
    Uptime: 00:03:49
    Downstream neighbors:
        Interface: Pseudo-GMP
            fe-0/0/0.0 fe-0/0/1.0 fe-0/0/3.0
        Interface: so-1/0/0.0 (pruned)
            10.111.10.2 State: Prune Flags: SR Timeout: 174
```

```
            Uptime: 00:03:49 Time since last Prune: 00:01:49
        Interface: mt-1/1/0.32768
            10.10.47.100 State: Join Flags: S   Timeout: Infinity
            Uptime: 00:03:49 Time since last Prune: 00:01:49
    Number of downstream interfaces: 3

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join extensive (Bidirectional PIM)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 233.252.0.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.13.2
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0
    Upstream neighbor: 10.10.1.2
    Upstream state: None
    Uptime: 00:03:49
    Bidirectional accepting interfaces:
        Interface: ge-0/0/1.0    (RPF)
        Interface: lo0.0         (DF Winner)
    Number of downstream interfaces: 0

Group: 233.252.0.1
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.13.2
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0
    Upstream neighbor: 10.10.1.2
    Upstream state: None
    Uptime: 00:03:49
    Bidirectional accepting interfaces:
        Interface: ge-0/0/1.0    (RPF)
        Interface: lo0.0         (DF Winner)
     Downstream neighbors:
         Interface: lt-1/0/10.24
             10.0.24.4 State: Join   RW  Timeout: 185
         Interface: lt-1/0/10.23
             10.0.23.3 State: Join   RW  Timeout: 184
     Number of downstream interfaces: 2

Group: 233.252.0.2
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.1.3
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0 (RP Link)
    Upstream neighbor: Direct
    Upstream state: Local RP
    Uptime: 00:03:49
    Bidirectional accepting interfaces:
        Interface: ge-0/0/1.0    (RPF)
        Interface: lo0.0         (DF Winner)
        Interface: xe-4/1/0.0    (DF Winner)
```

 Copyright © 2018, Juniper Networks, Inc.

```
    Number of downstream interfaces: 0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join extensive (Bidirectional PIM with a Directly Connected Phantom RP)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 233.252.0.0
    Bidirectional group prefix length: 24
    Source: *
    RP: 10.10.1.3
    Flags: bidirectional,rptree,wildcard
    Upstream interface: ge-0/0/1.0 (RP Link)
    Upstream neighbor: Direct
    Upstream state: Local RP
    Uptime: 00:03:49
    Bidirectional accepting interfaces:
        Interface: ge-0/0/1.0    (RPF)
        Interface: lo0.0         (DF Winner)
        Interface: xe-4/1/0.0    (DF Winner)
    Number of downstream interfaces: 0
```

show pim join instance <instance-name> extensive

```
user@host> show pim join instance VPN-A extensive
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 233.252.0.2
    Source: *
    RP: 10.10.47.100
    Flags: sparse,rptree,wildcard
    Upstream interface: Local
    Upstream neighbor: Local
    Upstream state: Local RP
    Uptime: 00:03:49
    Downstream neighbors:
        Interface: mt-1/1/0.32768
            10.10.47.101 State: Join Flags: SRW Timeout: 156
            Uptime: 00:03:49 Time since last Join: 00:01:49
    Number of downstream interfaces: 1

Group: 233.252.0.2
    Source: 192.168.195.74
    Flags: sparse,spt
    Upstream interface: at-0/3/1.0
    Upstream neighbor: 10.111.30.2
    Upstream state: Local RP, Join to Source
    Keepalive timeout: 156
    Uptime: 00:14:52

Group: 233.252.0.2
    Source: 192.168.195.169
    Flags: sparse
    Upstream interface: so-1/0/1.0
    Upstream neighbor: 10.111.20.2
```

Copyright © 2018, Juniper Networks, Inc.

```
    Upstream state: Local RP, Join to Source
    Keepalive timeout: 156
    Uptime: 00:14:52
```

show pim join extensive (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 233.252.0.1
    Source: 192.168.219.11
    Flags: sparse,spt
    Upstream interface: fe-1/3/1.0
    Upstream neighbor: Direct
    Upstream state: Local Source
    Keepalive timeout:
    Uptime: 11:27:55
    Downstream neighbors:
        Interface: Pseudo-MLDP
        Interface: lt-1/2/0.25
            10.2.5.2 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:27:55 Time since last Join: 11:27:55

Group: 233.252.0.2
    Source: 192.168.219.11
    Flags: sparse,spt
    Upstream interface: fe-1/3/1.0
    Upstream neighbor: Direct
    Upstream state: Local Source
    Keepalive timeout:
    Uptime: 11:27:41
    Downstream neighbors:
        Interface: Pseudo-MLDP

Group: 233.252.0.3
    Source: 192.168.219.11
    Flags: sparse,spt
    Upstream interface: fe-1/3/1.0
    Upstream neighbor: Direct
    Upstream state: Local Source
    Keepalive timeout:
    Uptime: 11:27:41
    Downstream neighbors:
        Interface: Pseudo-MLDP

Group: 233.252.0.22
    Source: 10.2.7.7
    Flags: sparse,spt
    Upstream interface: lt-1/2/0.27
    Upstream neighbor: Direct
    Upstream state: Local Source
    Keepalive timeout:
    Uptime: 11:27:25
    Downstream neighbors:
        Interface: Pseudo-MLDP

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 2001:db8::1:2
```

 Copyright © 2018, Juniper Networks, Inc.

```
    Source: 2001:db8::1:2:7:7
    Flags: sparse,spt
    Upstream interface: lt-1/2/0.27
    Upstream neighbor: Direct
    Upstream state: Local Source
    Keepalive timeout:
    Uptime: 11:27:26
    Downstream neighbors:
        Interface: Pseudo-MLDP
```

show pim join extensive (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 233.252.0.0
    Source: *
    RP: 10.1.1.1
    Flags: sparse,rptree,wildcard
    Upstream interface: Local
    Upstream neighbor: Local
    Upstream state: Local RP
    Uptime: 11:31:33
    Downstream neighbors:
        Interface: fe-1/3/0.0
            192.168.209.9 State: Join Flags: SRW  Timeout: Infinity
            Uptime: 11:31:33 Time since last Join: 11:31:32

Group: 233.252.0.1
    Source: 192.168.219.11
    Flags: sparse,spt
    Upstream protocol: MLDP
    Upstream interface: Pseudo MLDP
    Upstream neighbor: MLDP LSP root <10.1.1.2>
    Upstream state: Join to Source
    Keepalive timeout:
    Uptime: 11:31:32
    Downstream neighbors:
        Interface: so-0/1/3.0
            192.168.92.9 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:31:30 Time since last Join: 11:31:30
    Downstream neighbors:
        Interface: fe-1/3/0.0
            192.168.209.9 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:31:32 Time since last Join: 11:31:32

Group: 233.252.0.2
    Source: 192.168.219.11
    Flags: sparse,spt
    Upstream protocol: MLDP
    Upstream interface: Pseudo MLDP
    Upstream neighbor: MLDP LSP root <10.1.1.2>
    Upstream state: Join to Source
    Keepalive timeout:
    Uptime: 11:31:32
    Downstream neighbors:
        Interface: so-0/1/3.0
            192.168.92.9 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:31:30 Time since last Join: 11:31:30
    Downstream neighbors:
```

```
        Interface: lt-1/2/0.14
            10.1.4.4 State: Join Flags: S Timeout: 177
            Uptime: 11:30:33 Time since last Join: 00:00:33
    Downstream neighbors:
        Interface: fe-1/3/0.0
            192.168.209.9 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:31:32 Time since last Join: 11:31:32

Group: 233.252.0.3
    Source: 192.168.219.11
    Flags: sparse,spt
    Upstream protocol: MLDP
    Upstream interface: Pseudo MLDP
    Upstream neighbor: MLDP LSP root <10.1.1.2>
    Upstream state: Join to Source
    Keepalive timeout:
    Uptime: 11:31:32
    Downstream neighbors:
        Interface: fe-1/3/0.0
            192.168.209.9 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:31:32 Time since last Join: 11:31:32

Group: 233.252.0.22
    Source: 10.2.7.7
    Flags: sparse,spt
    Upstream protocol: MLDP
    Upstream interface: Pseudo MLDP
    Upstream neighbor: MLDP LSP root <10.1.1.2>
    Upstream state: Join to Source
    Keepalive timeout:
    Uptime: 11:31:30
    Downstream neighbors:
        Interface: so-0/1/3.0
            192.168.92.9 State: Join Flags: S   Timeout: Infinity
            Uptime: 11:31:30 Time since last Join: 11:31:30

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 2001:db8::1:2
    Source: 2001:db8::1:2:7:7
    Flags: sparse,spt
    Upstream protocol: MLDP
    Upstream interface: Pseudo MLDP
    Upstream neighbor: MLDP LSP root <10.1.1.2>
    Upstream state: Join to Source
    Keepalive timeout:
    Uptime: 11:31:32
    Downstream neighbors:
        Interface: fe-1/3/0.0
           2001:db8::21f:12ff:fea5:c4db State: Join Flags: S   Timeout: Infinity

            Uptime: 11:31:32 Time since last Join: 11:31:32
```

 Copyright © 2018, Juniper Networks, Inc.

show pim mdt

Syntax

show pim mdt instance *instance-name*
<brief | detail | extensive>
data-mdt-joins
data-mdt-limit
inet
inet6
<incoming | outgoing>
<logical-system (all | logical-system-name)>
<*range*>

Release Information

Command introduced before Junos OS Release 7.4.
Support for IPv6 added in Junos OS Release 17.3R1.

Description

Display information about Protocol Independent Multicast (PIM) default multicast distribution tree (MDT) and the data MDTs in a Layer 3 VPN environment for a routing instance.

Options

instance ***instance-name***—Display information about data-MDTs for a specific PIM-enabled routing instance.

brief | detail | extensive—(Optional) Display the specified level of output.

data-mdt-joins— Show received PIM data-mdt-joins.

data-mdt-limits— Show received PIM data-mdt-limits.

incoming | outgoing—(Optional) Display incoming or outgoing multicast data tunnels, respectively.

inet | inet6—Display IPv4 or IPv6 multicast data tunnels.

logical-system (all | ***logical-system-name*****)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

range—(Optional) Display information about an IP address with optional prefix length representing a particular multicast group.

Required Privilege Level

view

List of Sample Output

show pim mdt <variables> instance on page 1908
show pim mdt instance detail on page 1909
show pim mdt instance extensive on page 1909
show pim mdt instance incoming on page 1910
show pim mdt instance outgoing on page 1910
show pim mdt instance (SSM Mode) on page 1910

Output Fields Table 137 on page 1908 describes the output fields for the **show pim mdt** command. Output fields are listed in the approximate order in which they appear.

Table 137: show pim mdt Output Fields

Field Name	Field Description	Level of Output
Instance	Name of the routing instance.	All levels
Tunnel direction	Direction the tunnel faces, from the router's perspective: **Outgoing** or **Incoming**.	All levels
Tunnel mode	Mode the tunnel is operating in: **PIM-SSM** or **PIM-ASM**.	All levels
Default group address	Default multicast group address using this tunnel.	All levels
Default source address	Default multicast source address using this tunnel.	All levels
Default tunnel interface	Default multicast tunnel interface.	All levels
Default tunnel source	Address used as the source address for outgoing PIM control messages.	All levels
C-Group	Customer-facing multicast group address using this tunnel. If you enable dynamic reuse of data MDT group addresses, more than one group address can use the same data MDT.	**detail**
C-Source	IP address of the multicast source in the customer's address space. If you enable dynamic reuse of data MDT group addresses, more than one source address can use the same data MDT.	**detail**
P-Group	Service provider-facing multicast group address using this tunnel.	**detail**
Data tunnel interface	Multicast data tunnel interface that set up the data-MDT tunnel.	**detail**
Last known forwarding rate	Last known rate, in kilobits per second, at which the tunnel was forwarding traffic.	**detail**
Configured threshold rate	Rate, in kilobits per second, above which a data-MDT tunnel is created and below which it is deleted.	**detail**
Tunnel uptime	Time that this data-MDT tunnel has existed. The format is ***hours:minutes:seconds***.	**detail**

Sample Output

show pim mdt <variables> instance

Use this command to display MDT information for default MDT and data-MDT for IPv4 and/or IPv6 traffic.)

Copyright © 2018, Juniper Networks, Inc.

```
user@host> show pim mdt inet | inet6 instance VPN-A
Instance: PIM.VPN-A  Family: INET
Tunnel direction: Outgoing
Tunnel mode: PIM-SM
Default group address: 224.1.1.1
Default source address: 0.0.0.0
Default tunnel interface: mt-0/0/0.32768
Default tunnel source: 0.0.0.0

C-group address   C-source address   P-group address    Data tunnel interface
227.1.1.1         18.1.1.2           228.1.1.1          mt-0/0/0.32769

Instance: PIM.VPN-A
Tunnel direction: Incoming
Tunnel mode: PIM-SM
Default group address: 224.1.1.1
Default source address: 0.0.0.0
Default tunnel interface: mt-0/0/0.1081344
Default tunnel source: 0.0.0.0

Instance: PIM.VPN-A Family: INET6
```

show pim mdt instance detail

```
user@host> show pim mdt instance VPN-A detail
Instance: PIM.VPN-A
Tunnel direction: Outgoing
Default group address: 239.1.1.1
Default tunnel interface: mt-1/1/0.32768
Default tunnel source: 192.168.7.1

C-Group: 235.1.1.2
    C-Source: 192.168.195.74
    P-Group : 228.0.0.0
    Data tunnel interface      : mt-1/1/0.32769
    Last known forwarding rate : 48 kbps (6 kBps)
    Configured threshold rate  : 10 kbps
    Tunnel uptime              : 00:00:34

Instance: PIM.VPN-A
Tunnel direction: Incoming
Default group address: 239.1.1.1
Default tunnel interface: mt-1/1/0.1081344
```

show pim mdt instance extensive

```
user@host> show pim mdt instance VPN-A extensive
Instance: PIM.VPN-A
Tunnel direction: Outgoing
Default group address: 239.1.1.1
Default tunnel interface: mt-1/1/0.32768
Default tunnel source: 192.168.7.1

C-Group: 235.1.1.2
    C-Source: 192.168.195.74
    P-Group : 228.0.0.0
    Data tunnel interface      : mt-1/1/0.32769
    Last known forwarding rate : 48 kbps (6 kBps)
    Configured threshold rate  : 10 kbps
    Tunnel uptime              : 00:00:41
```

```
Instance: PIM.VPN-A
Tunnel direction: Incoming
Default group address: 239.1.1.1
Default tunnel interface: mt-1/1/0.1081344
```

show pim mdt instance incoming

```
user@host> show pim mdt instance VPN-A incoming
Instance: PIM.VPN-A
Tunnel direction: Incoming
Default group address: 239.1.1.1
Default tunnel interface: mt-1/1/0.1081344
```

show pim mdt instance outgoing

```
user@host> show pim mdt instance VPN-A outgoing
Instance: PIM.VPN-A
Tunnel direction: Outgoing
Default group address: 239.1.1.1
Default tunnel interface: mt-1/1/0.32768
Default tunnel source: 192.168.7.1

C-group address   C-source address   P-group address    Data tunnel interface
235.1.1.2         192.168.195.74     228.0.0.0          mt-1/1/0.32769
```

show pim mdt instance (SSM Mode)

```
user@host> show pim mdt instance vpn-a
Instance: PIM.vpn-a
Tunnel direction: Outgoing
Tunnel mode: PIM-SSM
Default group address: 232.1.1.1
Default source address: 10.255.14.216
Default tunnel interface: mt-1/3/0.32769
Default tunnel source: 192.168.7.1

Instance: PIM.vpn-a
Tunnel direction: Incoming
Tunnel mode: PIM-SSM
Default group address: 232.1.1.1
Default source address: 10.255.14.217
Default tunnel interface: mt-1/3/0.1081345

Instance: PIM.vpn-a
Tunnel direction: Incoming
Tunnel mode: PIM-SSM
Default group address: 232.1.1.1
Default source address: 10.255.14.218
Default tunnel interface: mt-1/3/0.1081345
```

 Copyright © 2018, Juniper Networks, Inc.

show pim mdt data-mdt-joins

Syntax

show pim mdt data-mdt-joins
<logical-system (all | *logical-system-name*)> instance *instance-name*

Release Information

Command introduced in Junos OS Release 11.2.

Description

In a draft-rosen Layer 3 multicast virtual private network (MVPN) configured with service provider tunnels, display the advertisements of new multicast distribution tree (MDT) group addresses cached by the provider edge (PE) routers in the specified VPN routing and forwarding (VRF) instance that is configured to use the Protocol Independent Multicast (PIM) protocol.

Options

instance *instance-name*—Display data MDT join packets cached by PE routers in a specific PIM instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

NOTE: **Draft-rosen multicast VPNs are not supported in a logical system environment even though the configuration statements can be configured under the logical-systems hierarchy.**

Required Privilege Level

view

Related Documentation

- *Understanding Data MDTs*
- *Example: Configuring Data MDTs and Provider Tunnels Operating in Source-Specific Multicast Mode*
- *Example: Configuring Data MDTs and Provider Tunnels Operating in Any-Source Multicast Mode*

List of Sample Output

show pim mdt data-mdt-joins on page 1912

Output Fields

Table 138 on page 1912 describes the output fields for the **show pim mdt data-mdt-joins** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2018, Juniper Networks, Inc.

Table 138: show pim mdt data-mdt-joins Output Fields

Field Name	Field Description
C-Group	IPv4 group address in the address space of the customer's VPN-specific PIM-enabled routing instance of the multicast traffic destination. This 32-bit value is carried in the C-group field of the MDT join TLV packet.
C-Source	IPv4 address in the address space of the customer's VPN-specific PIM-enabled routing instance of the multicast traffic source. This 32-bit value is carried in the C-source field of the MDT join TLV packet.
P-Group	IPv4 group address in the service provider's address space of the new data MDT that the PE router will use to encapsulate the VPN multicast traffic flow (C-Source, C-Group). This 32-bit value is carried in the P-group field of the MDT join TLV packet.
P-Source	IPv4 address of the PE router.
Timeout	Timeout, in seconds, remaining for this cache entry. When the cache entry is created, this field is set to 180 seconds. After an entry times out, the PE router deletes the entry from its cache and prunes itself off the data MDT.

Sample Output

show pim mdt data-mdt-joins

```
user@host show pim mdt data-mdt-joins instance VPN-A
C-Source        C-Group         P-Source        P-Group         Timeout
20.2.15.9       225.1.1.2       20.0.0.5        239.10.10.0     172
20.2.15.9       225.1.1.3       20.0.0.5        239.10.10.1     172
```

 Copyright © 2018, Juniper Networks, Inc.

show pim mdt data-mdt-limit

Syntax

show pim mdt data-mdt-limit instance *instance-name*
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced in Junos OS Release 12.2.

Description

Display the maximum number configured and the currently active data multicast distribution trees (MDTs) for a specific VPN routing and forwarding (VRF) instance.

Options

instance ***instance-name***—Display data MDT information for the specified VRF instance.

logical-system (all | ***logical-system-name*****)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

NOTE: Draft-rosen multicast VPNs are not supported in a logical system environment even though the configuration statements can be configured under the logical-systems hierarchy.

Required Privilege Level

view

Related Documentation

- *Understanding Data MDTs*
- *Example: Configuring Data MDTs and Provider Tunnels Operating in Source-Specific Multicast Mode*
- *Example: Configuring Data MDTs and Provider Tunnels Operating in Any-Source Multicast Mode*

List of Sample Output

show pim mdt data-mdt-limit on page 1914

Output Fields

Table 139 on page 1913 describes the output fields for the **show pim mdt data-mdt-limit** command. Output fields are listed in the approximate order in which they appear.

Table 139: show pim mdt data-mdt-limit Output Fields

Field Name	Field Description
Maximum Data Tunnels	Maximum number of data MDTs created in this VRF instance. If the number is 0, no data MDTs are created for this VRF instance.
Active Data Tunnels	Active number of data MDTs in this VRF instance.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show pim mdt data-mdt-limit

```
user@host show pim mdt data-mdt-limit instance VPN-A
Maximum Data Tunnels                 10
Active Data Tunnels                   2
```


Copyright © 2018, Juniper Networks, Inc.

show pim neighbors

List of Syntax

Syntax on page 1915
Syntax (EX Series Switch and the QFX Series) on page 1915

Syntax

show pim neighbors
<brief | detail>
<inet | inet6>
<instance (*instance-name* | all)>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)

show pim neighbors
<brief | detail>
<inet | inet6>
<instance (*instance-name* | all)>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Support for bidirectional PIM added in Junos OS Release 12.1.
Support for the **instance all** option added in Junos OS Release 12.1.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Support for RFC 5496 (via **rpf-vector**) added in Junos OS Release 17.3R1.

Description

Display information about Protocol Independent Multicast (PIM) neighbors.

Options

none—(Same as **brief**) Display standard information about PIM neighbors for all supported family addresses for the main instance.

brief | detail—(Optional) Display the specified level of output.

inet | inet6—(Optional) Display information about PIM neighbors for IPv4 or IPv6 family addresses, respectively.

instance (*instance-name* | all)—(Optional) Display information about neighbors for the specified PIM-enabled routing instance or for all routing instances.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show pim neighbors on page 1917
show pim neighbors instance on page 1917
show pim neighbors detail on page 1917
show pim neighbors detail (With BFD) on page 1918

Copyright © 2018, Juniper Networks, Inc.

Output Fields Table 140 on page 1916 describes the output fields for the **show pim neighbors** command. Output fields are listed in the approximate order in which they appear.

Table 140: show pim neighbors Output Fields

Field Name	Field Description	Level of Output
Instance	Name of the routing instance.	All levels
Interface	Interface through which the neighbor is reachable.	All levels
Neighbor addr	Address of the neighboring PIM routing device.	All levels
IP	IP version: **4** or **6**.	All levels
V	PIM version running on the neighbor: **1** or **2**.	All levels
Mode	PIM mode of the neighbor: **Sparse**, **Dense**, **SparseDense**, or **Unknown**. When the neighbor is running PIM version 2, this mode is always **Unknown**.	All levels
Option	Can be one or more of the following: • **B**—Bidirectional Capable. • **G**—Generation Identifier. • **H**—Hello Option Holdtime. • **L**—Hello Option LAN Prune Delay. • **P**—Hello Option DR Priority. • **T**—Tracking bit. • **A**—Join attribute; used in conjunction with **pim rpf-vector**.	**brief** none
Uptime	Time the neighbor has been operational since the PIM process was last initialized. Starting in Junos OS release 17.3R1, **uptime** is not reset during ISSU.The time format is as follows: ***dd*:*hh*:*mm*:*ss* ago** for less than a week and ***nwnd*:*hh*:*mm*:*ss* ago** for more than a week.	All levels
Address	Address of the neighboring PIM routing device.	**detail**
BFD	Status and operational state of the Bidirectional Forwarding Detection (BFD) protocol on the interface: **Enabled**, **Operational state is up**, or **Disabled**.	**detail**
Hello Option Holdtime	Time for which the neighbor is available, in seconds. The range of values is **0** through **65,535.**	**detail**
Hello Default Holdtime	Default holdtime and the time remaining if the **holdtime** option is not in the received hello message.	**detail**
Hello Option DR Priority	Designated router election priority. The range of values is **0** through **255**.	**detail**
Hello Option Join Attribute	Appears in conjunction with the **rpf-vector** command. The Join attribute is included in the PIM join messages of PIM routers that can receive type 1 Encoded-Source Address.	**detail**

 Copyright © 2018, Juniper Networks, Inc.

Table 140: show pim neighbors Output Fields (continued)

Field Name	Field Description	Level of Output
Hello Option Generation ID	9-digit or 10-digit number used to tag hello messages.	**detail**
Hello Option Bi-Directional PIM supported	Neighbor can process bidirectional PIM messages.	**detail**
Hello Option LAN Prune Delay	Time to wait before the neighbor receives prune messages, in the format **delay** ***nnn*** **ms override** ***nnnn*** **ms**.	**detail**
Join Suppression supported	Neighbor is capable of join suppression.	**detail**
Rx Join	Information about joins received from the neighbor. • **Group**—Group addresses in the join message. • **Source**—Address of the source in the join message. • **Timeout**—Time for which the join is valid.	**detail**

Sample Output

show pim neighbors

```
user@host> show pim neighbors
Instance: PIM.master
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking bit
A = Hello Option Join Attribute

Instance: PIM.master
Interface  IP V Mode        Option       Uptime Neighbor addr
ae0.0      4  2             HPLGTA       19:01:24 20.0.0.13
ae1.0      4  2             HPLGTA       19:01:24 20.0.0.149
```

show pim neighbors instance

```
user@host> show pim neighbors instance VPN-A
Instance: PIM.VPN-A
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking bit

Interface           IP V Mode        Option      Uptime Neighbor addr
at-0/3/1.0           4 2             HPLG      00:07:54 10.111.30.2
mt-1/1/0.32768       4 2             HPLG      00:07:22 10.10.47.101
so-1/0/1.0           4 2             HPLG      00:07:50 10.111.20.2
```

show pim neighbors detail

```
user@host> show pim neighbors detail
```

Copyright © 2018, Juniper Networks, Inc.

```
Instance: PIM.master
Interface: ae1.0

Address: 20.0.0.149, IPv4, PIM v2, sg Join Count: 0, tsg Join Count: 332
        BFD: Disabled
        Hello Option Holdtime: 105 seconds 86 remaining
        Hello Option DR Priority: 1
        Hello Option Generation ID: 853386212
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
                                 Join Suppression supported
        Hello Option Join Attribute supported

    Address: 20.0.0.150, IPv4, PIM v2, Mode: SparseDense, sg Join Count: 0, tsg
Join Count: 0
        Hello Option Holdtime: 65535 seconds
        Hello Option DR Priority: 1
        Hello Option Generation ID: 358917871
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
                                 Join Suppression supported
             Hello Option Join Attribute supported

Interface: lo0.0

    Address: 10.255.179.246,    IPv4, PIM v2, Mode: SparseDense, sg Join Count:
0, tsg Join Count: 0
        Hello Option Holdtime: 65535 seconds
        Hello Option DR Priority: 1
        Hello Option Generation ID: 1997462267
        Hello Option Bi-Directional PIM supported
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
                                      Join Suppression supported
```

show pim neighbors detail (With BFD)

```
user@host> show pim neighbors detail
Instance: PIM.master
Interface: fe-1/0/0.0
    Address: 192.168.11.1,        IPv4, PIM v2, Mode: Sparse
        Hello Option Holdtime: 65535 seconds
        Hello Option DR Priority: 1
        Hello Option Generation ID: 836607909
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

    Address: 192.168.11.2,        IPv4, PIM v2
        BFD: Enabled, Operational state is up
        Hello Default Holdtime: 105 seconds 104 remaining
        Hello Option DR Priority: 1
        Hello Option Generation ID: 1907549685
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

Interface: fe-1/0/1.0
    Address: 192.168.12.1,        IPv4, PIM v2
        BFD: Disabled
        Hello Default Holdtime: 105 seconds 80 remaining
        Hello Option DR Priority: 1
        Hello Option Generation ID: 1971554705
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

 Copyright © 2018, Juniper Networks, Inc.

show pim rps

List of Syntax

Syntax on page 1919
Syntax (EX Series Switch and the QFX Series) on page 1919

Syntax

show pim rps
<brief | detail | extensive>
<*group-address*>
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)

show pim rps
<brief | detail | extensive>
<*group-address*>
<inet | inet6>
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Support for bidirectional PIM added in Junos OS Release 12.1.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display information about Protocol Independent Multicast (PIM) rendezvous points (RPs).

Options

none—Display standard information about PIM RPs for all groups and family addresses for all routing instances.

brief | detail | extensive—(Optional) Display the specified level of output.

group-address—(Optional) Display the RPs for a particular group. If you specify a group address, the output lists the routing device that is the RP for that group.

inet | inet6—(Optional) Display information for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display information about RPs for a specific PIM-enabled routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- *Example: Configuring Bidirectional PIM*

Copyright © 2018, Juniper Networks, Inc.

List of Sample Output

Output Fields

Table 141 on page 1920 describes the output fields for the **show pim rps** command. Output fields are listed in the approximate order in which they appear.

Table 141: show pim rps Output Fields

Field Name	Field Description	Level of Output
Instance	Name of the routing instance.	All levels
Family or Address family	Name of the address family: **inet** (IPv4) or **inet6** (IPv6).	All levels
RP address	Address of the rendezvous point.	All levels
Type	Type of RP: • **auto-rp**—Address of the RP known through the Auto-RP protocol. • **bootstrap**—Address of the RP known through the bootstrap router protocol (BSR). • **embedded**—Address of the RP known through an embedded RP (IPv6). • **static**—Address of RP known through static configuration.	**brief** none
Holdtime	How long to keep the RP active, with time remaining, in seconds.	All levels
Timeout	How long until the local routing device determines the RP to be unreachable, in seconds.	All levels
Groups	Number of groups currently using this RP.	All levels
Group prefixes	Addresses of groups that this RP can span.	**brief** none
Learned via	Address and method by which the RP was learned.	**detail extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 141: show pim rps Output Fields (continued)

Field Name	Field Description	Level of Output
Mode	The PIM mode of the RP: bidirectional or sparse. If a sparse and bidirectional RPs are configured with the same RP address, they appear as separate entries in both formats.	All levels
Time Active	How long the RP has been active, in the format ***hh:mm:ss***.	**detail extensive**
Device Index	Index value of the order in which Junos OS finds and initializes the interface. For bidirectional RPs, the **Device Index** output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.	**detail extensive**
Subunit	Logical unit number of the interface. For bidirectional RPs, the **Subunit** output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.	**detail extensive**
Interface	Either the encapsulation or the de-encapsulation logical interface, depending on whether this routing device is a designated router (DR) facing an RP router, or is the local RP, respectively. For bidirectional RPs, the **Interface** output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.	**detail extensive**
Group Ranges	Addresses of groups that this RP spans.	**detail extensive** *group-address*
Active groups using RP	Number of groups currently using this RP.	**detail extensive**
total	Total number of active groups for this RP.	**detail extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 141: show pim rps Output Fields (continued)

Field Name	Field Description	Level of Output
Register State for RP	Current register state for each group: • **Group**—Multicast group address. • **Source**—Multicast source address for which the PIM register is sent or received, depending on whether this router is a designated router facing an RP router, or is the local RP, respectively: • **First Hop**—PIM-designated routing device that sent the Register message (the source address in the IP header). • **RP Address**—RP to which the Register message was sent (the destination address in the IP header). • **State**: On the designated router: ◦ **Send**—Sending Register messages. ◦ **Probe**—Sent a null register. If a Register-Stop message does not arrive in 5 seconds, the designated router resumes sending Register messages. ◦ **Suppress**—Received a Register-Stop message. The designated router is waiting for the timer to resume before changing to **Probe** state. • On the RP: ◦ **Receive**—Receiving Register messages.	**extensive**
Anycast-PIM rpset	If anycast RP is configured, the addresses of the RPs in the set.	**extensive**
Anycast-PIM local address used	If anycast RP is configured, the local address used by the RP.	**extensive**
Anycast-PIM Register State	If anycast RP is configured, the current register state for each group: • **Group**—Multicast group address. • **Source**—Multicast source address for which the PIM register is sent or received, depending on whether this routing device is a designated router facing an RP router, or is the local RP, respectively. • **Origin**—How the information was obtained: ◦ **DIRECT**—From a local attachment ◦ **MSDP**—From the Multicast Source Discovery Protocol (MSDP) ◦ **DR**—From the designated router	**extensive**
RP selected	For sparse mode and bidirectional mode, the identity of the RP for the specified group address.	*group-address*

Sample Output

show pim rps

```
user@host> show pim rps
Instance: PIM.master

Address-family INET
RP address       Type        Mode   Holdtime Timeout Groups Group prefixes
```


Copyright © 2018, Juniper Networks, Inc.

```
10.100.100.100 auto-rp     sparse      150      146       0 233.252.0.0/8
                                                            233.252.0.1/24
10.200.200.200 auto-rp     sparse      150      146       0 233.252.0.2/4

address-family INET6
```

show pim rps brief

The output for the **show pim rps brief** command is identical to that for the **show pim rps** command. For sample output, see show pim rps on page 1922.

show pim rps <group-address>

```
user@host> show pim rps 233.252.0.0
Instance: PIM.master
Instance: PIM.master

RP selected: 10.100.100.100
```

show pim rps <group-address>

```
user@host> show pim rps 233.252.0.0
Instance: PIM.master
Instance: PIM.master

RP selected: 10.100.100.100
```

show pim rps <group-address> (Bidirectional PIM)

```
user@host> show pim rps 233.252.0.1
Instance: PIM.master

233.252.0.0/16
        10.4.12.75 (Bidirectional)

RP selected: 10.4.12.75
```

show pim rps <group-address> (PIM Dense Mode)

```
user@host> show pim rps 233.252.0.1
Instance: PIM.master

Dense Mode active for group 233.252.0.1
```

show pim rps <group-address> (SSM Range Without asm-override-ssm Configured)

```
user@host> show pim rps 233.252.0.1
Instance: PIM.master

Source-specific Mode (SSM) active for group 233.252.0.1
```

show pim rps <group-address> (SSM Range With asm-override-ssm Configured and a Sparse-Mode RP)

```
user@host> show pim rps 233.252.0.1
Instance: PIM.master

Source-specific Mode (SSM) active with Sparse Mode ASM override for group
233.252.0.1
```

Copyright © 2018, Juniper Networks, Inc.

```
233.252.0.0/16
        10.4.12.75

RP selected: 10.4.12.75
```

show pim rps <group-address> (SSM Range With asm-override-ssm Configured and a Bidirectional RP)

```
user@host> show pim rps 233.252.0.1
Instance: PIM.master

Source-specific Mode (SSM) active with Sparse Mode ASM override for group
233.252.0.1

233.252.0.0/16
        10.4.12.75 (Bidirectional)

RP selected: (null)
```

show pim rps instance

```
user@host> show pim rps instance VPN-A
Instance: PIM.VPN-A
Address family INET
RP address                Type        Holdtime Timeout Groups Group prefixes
10.10.47.100              static             0     None      1 233.252.0.0/4

Address family INET6
```

show pim rps extensive (PIM Sparse Mode)

```
user@host> show pim rps extensive
Instance: PIM.master

Family: INET
RP: 10.255.245.91
Learned via: static configuration
Time Active: 00:05:48
Holdtime: 45 with 36 remaining
Device Index: 122
Subunit: 32768
Interface: pd-6/0/0.32768
Group Ranges:
    233.252.0.0/4, 36s remaining
Active groups using RP:
    233.252.0.1

    total 1 groups active

Register State for RP:
Group           Source          FirstHop        RP Address      State    Timeout
233.252.0.1       192.168.195.78  10.255.14.132   10.255.245.91   Receive
 0
```

show pim rps extensive (Bidirectional PIM)

```
user@host> show pim rps extensive
Instance: PIM.master
Address family INET
```

 Copyright © 2018, Juniper Networks, Inc.

```
RP: 10.10.1.3
Learned via: static configuration
Mode: Bidirectional
Time Active: 01:58:07
Holdtime: 150
Group Ranges:
        233.252.0.0/24
        233.252.0.01/24

RP: 10.10.13.2
Learned via: static configuration
Mode: Bidirectional
Time Active: 01:58:07
Holdtime: 150
Group Ranges:
       233.252.0.3/24
       233.252.0.4/24
```

show pim rps extensive (PIM Anycast RP in Use)

```
user@host> show pim rps extensive
Instance: PIM.master

Family: INET
RP: 10.10.10.2
Learned via: static configuration
Time Active: 00:54:52
Holdtime: 0
Device Index: 130
Subunit: 32769
Interface: pimd.32769
Group Ranges:
    233.252.0.0/4
Active groups using RP:
    233.252.0.10

    total 1 groups active

Anycast-PIM rpset:
           10.100.111.34
           10.100.111.17
           10.100.111.55

Anycast-PIM local address used: 10.100.111.1
Anycast-PIM Register State:
Group                    Source                              Origin
233.252.0.1              10.10.95.2                          DIRECT
233.252.0.2              10.10.95.2                          DIRECT
233.252.0.3              10.10.70.1                          MSDP
233.252.0.4              10.10.70.1                          MSDP
233.252.0.5              10.10.71.1                          DR

Address family INET6

Anycast-PIM rpset:
           ab::1
           ab::2
Anycast-PIM local address used: cd::1

Anycast-PIM Register State:
Group                  Source                              Origin
```

```
::224.1.1.1            ::10.10.95.2                 DIRECT
::224.1.1.2            ::10.10.95.2                 DIRECT
::224.20.20.1          ::10.10.71.1                 DR
```

 Copyright © 2018, Juniper Networks, Inc.

show pim snooping interfaces

Syntax

show pim snooping interfaces
<brief | detail>
<instance *instance-name*>
<interface *interface-name*>
<logical-system *logical-system-name*>
<vlan-id *vlan-identifier*>

Release Information

Command introduced in Junos OS Release 12.3 for MX Series 3D Universal Edge devices.
Command introduced in Junos OS Release 13.2 for M Series Multiservice Edge devices.

Description

Display information about PIM snooping interfaces.

Options

none—Display detailed information.

brief | detail—(Optional) Display the specified level of output.

instance *<instance-name>*—(Optional) Display PIM snooping interface information for the specified routing instance.

interface *<interface-name>*—(Optional) Display PIM snooping information for the specified interface only.

logical-system *logical-system-name*—(Optional) Display information about a particular logical system, or type 'all'.

vlan-id *<vlan-identifier>*—(Optional) Display PIM snooping interface information for the specified VLAN.

Required Privilege Level

view

Related Documentation

- *PIM Snooping for VPLS*

List of Sample Output

show pim snooping interfaces on page 1928
show pim snooping interfaces instance vpls1 on page 1928
show pim snooping interfaces interface <interface-name> on page 1929
show pim snooping interfaces vlan-id <vlan-id> on page 1929

Output Fields

Table 142 on page 1927 lists the output fields for the **show pim snooping interface** command. Output fields are listed in the approximate order in which they appear.

Table 142: show pim snooping interface Output Fields

Field Name	Field Description	Level of Output
Instance	Routing instance for PIM snooping.	All levels

Copyright © 2018, Juniper Networks, Inc.

Table 142: show pim snooping interface Output Fields (continued)

Field Name	Field Description	Level of Output
Learning-Domain	Learning domain for snooping.	All levels
Name	Router interfaces that are part of this learning domain.	All levels
State	State of the interface: **Up**, or **Down**.	All levels
IP-Version	Version of IP used: **4** for IPv4, or **6** for IPv6.	All levels
NbrCnt	Number of neighboring routers connected through the specified interface.	All levels
DR address	IP address of the designated router.	All levels

Sample Output

show pim snooping interfaces

```
user@host> show pim snooping interfaces
Instance: vpls1
Learning-Domain: vlan-id 10
Name State IP-Version NbrCnt
ge-1/3/1.10 Up 4 1
ge-1/3/3.10 Up 4 1
ge-1/3/5.10 Up 4 1
ge-1/3/7.10 Up 4 1
DR address: 192.0.2.5
DR flooding is ON

Learning-Domain: vlan-id 20
Name State IP-Version NbrCnt
ge-1/3/1.20 Up 4 1
ge-1/3/3.20 Up 4 1
ge-1/3/5.20 Up 4 1
ge-1/3/7.20 Up 4 1
DR address: 192.0.2.6
DR flooding is ON
```

show pim snooping interfaces instance vpls1

```
user@host> show pim snooping interfaces instance vpls1
Instance: vpls1

Learning-Domain: vlan-id 10
Name State IP-Version NbrCnt
ge-1/3/1.10 Up 4 1
ge-1/3/3.10 Up 4 1
ge-1/3/5.10 Up 4 1
ge-1/3/7.10 Up 4 1
DR address: 192.0.2.5
DR flooding is ON

Learning-Domain: vlan-id 20
Name State IP-Version NbrCnt
```

 Copyright © 2018, Juniper Networks, Inc.

```
ge-1/3/1.20 Up 4 1
ge-1/3/3.20 Up 4 1
ge-1/3/5.20 Up 4 1
ge-1/3/7.20 Up 4 1
DR address: 192.0.2.6
DR flooding is ON
```

show pim snooping interfaces interface <interface-name>

```
user@host> show pim snooping interfaces interface ge-1/3/1.10
Instance: vpls1
Learning-Domain: vlan-id 10

Name State IP-Version NbrCnt
ge-1/3/1.10 Up 4 1
DR address: 192.0.2.5
DR flooding is ON

Learning-Domain: vlan-id 20
DR address: 192.0.2.6
DR flooding is ON
```

show pim snooping interfaces vlan-id <vlan-id>

```
user@host> show pim snooping interfaces vlan-id 10
Instance: vpls1
Learning-Domain: vlan-id 10

Name State IP-Version NbrCnt
ge-1/3/1.10 Up 4 1
ge-1/3/3.10 Up 4 1
ge-1/3/5.10 Up 4 1
ge-1/3/7.10 Up 4 1
DR address: 192.0.2.5
DR flooding is ON
```

Copyright © 2018, Juniper Networks, Inc.

show pim snooping join

Syntax

show pim snooping join
<brief | detail | extensive>
<instance *instance-name*>
<logical-system *logical-system-name*>
<vlan-id *vlan-id*>

Release Information

Command introduced in Junos OS Release 12.3 for MX Series 3D Universal Edge devices.
Command introduced in Junos OS Release 13.2 for M Series Multiservice Edge devices.

Description

Display information about Protocol Independent Multicast (PIM) snooping joins.

Options

none—Display detailed information.

brief | detail | extensive—(Optional) Display the specified level of output.

instance *instance-name*—(Optional) Display PIM snooping join information for the specified routing instance.

logical-system *logical-system-name*—(Optional) Display information about a particular logical system, or type 'all'.

vlan-id *vlan-identifier*—(Optional) Display PIM snooping join information for the specified VLAN.

Required Privilege Level

view

Related Documentation

- *PIM Snooping for VPLS*

List of Sample Output

show pim snooping join on page 1932
show pim snooping join extensive on page 1932
show pim snooping join instance on page 1932
show pim snooping join vlan-id on page 1933

Output Fields

Table 143 on page 1930 lists the output fields for the **show pim snooping join** command. Output fields are listed in the approximate order in which they appear.

Table 143: show pim snooping join Output Fields

Field Name	Field Description	Level of Output
Instance	Routing instance for PIM snooping.	All levels
Learning-Domain	Learning domain for PIM snooping.	All levels
Group	Multicast group address.	All levels

Copyright © 2018, Juniper Networks, Inc.

Table 143: show pim snooping join Output Fields (continued)

Field Name	Field Description	Level of Output
Source	Multicast source address: • * (wildcard value) • *<ipv4-address>* • *<ipv6-address>*	All levels
Flags	PIM flags: • **bidirectional**—Bidirectional mode entry. • **dense**—Dense mode entry. • **rptree**—Entry is on the rendezvous point tree. • **sparse**—Sparse mode entry. • **spt**—Entry is on the shortest-path tree for the source. • **wildcard**—Entry is on the shared tree.	All levels
Upstream state	Information about the upstream interface: • **Join to RP**—Sending a join to the rendezvous point. • **Join to Source**—Sending a join to the source. • **Local RP**—Sending neither join messages nor prune messages toward the RP, because this router is the rendezvous point. • **Local Source**—Sending neither join messages nor prune messages toward the source, because the source is locally attached to this routing device. • **Prune to RP**—Sending a prune to the rendezvous point. • **Prune to Source**—Sending a prune to the source. NOTE: RP group range entries have **None** in the **Upstream state** field because RP group ranges do not trigger actual PIM join messages between routers.	All levels
Upstream neighbor	Information about the upstream neighbor: **Direct**, **Local**, **Unknown**, or a specific IP address. For bidirectional PIM, **Direct** means that the interface is directly connected to a subnet that contains a phantom RP address.	All levels
Upstream port	RPF interface toward the source address for the source-specific state (S,G) or toward the rendezvous point (RP) address for the non-source-specific state (*,G). For bidirectional PIM, **RP Link** means that the interface is directly connected to a subnet that contains a phantom RP address.	All levels
Downstream port	Information about downstream interfaces.	**extensive**
Downstream neighbors	Address of the downstream neighbor.	**extensive**
Timeout	Time remaining until the downstream join state is updated (in seconds).	**extensive**

Sample Output

show pim snooping join

```
user@host> show pim snooping join
Instance: vpls1

Learning-Domain: vlan-id 10
Group: 198.51.100.2
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 192.0.2.4, port: ge-1/3/5.10


Learning-Domain: vlan-id 20
Group: 198.51.100.3
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 203.0.113.4, port: ge-1/3/5.20
```

show pim snooping join extensive

```
user@host> show pim snooping join extensive
Instance: vpls1
Learning-Domain: vlan-id 10

Group: 198.51.100.2
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 192.0.2.4, port: ge-1/3/5.10
Downstream port: ge-1/3/1.10
Downstream neighbors:
192.0.2.2 State: Join Flags: SRW Timeout: 166

Learning-Domain: vlan-id 20
Group: 198.51.100.3
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 203.0.113.4, port: ge-1/3/5.20
Downstream port: ge-1/3/3.20
Downstream neighbors:
203.0.113.3 State: Join Flags: SRW Timeout: 168
```

show pim snooping join instance

```
user@host> show pim snooping join instance vpls1
Instance: vpls1

Learning-Domain: vlan-id 10
Group: 198.51.100.2
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 192.0.2.4, port: ge-1/3/5.10
```

 Copyright © 2018, Juniper Networks, Inc.

```
Learning-Domain: vlan-id 20
Group: 198.51.100.3
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 203.0.113.4, port: ge-1/3/5.20
```

show pim snooping join vlan-id

```
user@host> show pim snooping join vlan-id 10
Instance: vpls1
Learning-Domain: vlan-id 10
Group: 198.51.100.2
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 192.0.2.4, port: ge-1/3/5.10
```

show pim snooping neighbors

Syntax

show pim snooping neighbors
<brief | detail>
<instance *instance-name*>
<interface *interface-name*>
<logical-system *logical-system-name*>
<vlan-id *vlan-identifier*>

Release Information

Command introduced in Junos OS Release 12.3 for MX Series 3D Universal Edge devices.
Command introduced in Junos OS Release 13.2 for M Series Multiservice Edge devices.

Description

Display information about Protocol Independent Multicast (PIM) snooping neighbors.

Options

none—Display detailed information.

brief | detail—(Optional) Display the specified level of output.

instance *instance-name*—(Optional) Display PIM snooping neighbor information for the specified routing instance.

interface *interface-name*—(Optional) Display information for the specified PIM snooping neighbor interface.

logical-system *logical-system-name*—(Optional) Display information about a particular logical system, or type 'all'.

vlan-id *vlan-identifier*—(Optional) Display PIM snooping neighbor information for the specified VLAN.

Required Privilege Level

view

Related Documentation

- *Configuring Interface Priority for PIM Designated Router Selection*
- *Modifying the PIM Hello Interval*
- *PIM Snooping for VPLS*
- show pim neighbors on page 1915

List of Sample Output

show pim snooping neighbors on page 1935
show pim snooping neighbors detail on page 1936
show pim snooping neighbors instance on page 1937
show pim snooping neighbors interface on page 1937
show pim snooping neighbors vlan-id on page 1938

Output Fields

Table 144 on page 1935 lists the output fields for the **show pim snooping neighbors** command. Output fields are listed in the approximate order in which they appear.

 Copyright © 2018, Juniper Networks, Inc.

Table 144: show pim snooping neighbors Output Fields

Field Name	Field Description	Level of Output
Instance	Routing instance for PIM snooping.	All levels
Learning-Domain	Learning domain for PIM snooping.	All levels
Interface	Router interface for which PIM snooping neighbor details are displayed.	All levels
Option	PIM snooping options available on the specified interface: • H = Hello Option Holdtime • P = Hello Option DR Priority • L = Hello Option LAN Prune Delay • G = Generation Identifier • T = Tracking Bit	All levels
Uptime	Time the neighbor has been operational since the PIM process was last initialized, in the format ***dd*:*hh*:*mm*:*ss* ago** for less than a week and ***nwnd*:*hh*:*mm*:*ss* ago** for more than a week.	All levels
Neighbor addr	IP address of the PIM snooping neighbor connected through the specified interface.	All levels
Address	IP address of the specified router interface.	All levels
Hello Option Holdtime	Time for which the neighbor is available, in seconds. The range of values is **0** through **65,535**.	**detail**
Hello Option DR Priority	Designated router election priority. The range of values is **0** through **4294967295**. NOTE: By default, every PIM interface has an equal probability (priority 1) of being selected as the DR.	**detail**
Hello Option Generation ID	9-digit or 10-digit number used to tag hello messages.	**detail**
Hello Option LAN Prune Delay	Time to wait before the neighbor receives prune messages, in the format **delay *nnn* ms override *nnnn* ms**.	**detail**

Sample Output

show pim snooping neighbors

```
user@host> show pim snooping neighbors
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking Bit

Instance: vpls1
Learning-Domain: vlan-id 10

Interface Option Uptime Neighbor addr
ge-1/3/1.10 HPLGT 00:43:33 192.0.2.2
```

Copyright © 2018, Juniper Networks, Inc.

```
ge-1/3/3.10 HPLGT 00:43:33 192.0.2.3
ge-1/3/5.10 HPLGT 00:43:33 192.0.2.4
ge-1/3/7.10 HPLGT 00:43:33 192.0.2.5

Learning-Domain: vlan-id 20

Interface Option Uptime Neighbor addr
ge-1/3/1.20 HPLGT 00:43:33 192.0.2.12
ge-1/3/3.20 HPLGT 00:43:33 192.0.2.13
ge-1/3/5.20 HPLGT 00:43:33 192.0.2.14
ge-1/3/7.20 HPLGT 00:43:33 192.0.2.15
```

show pim snooping neighbors detail

```
user@host> show pim snooping neighbors detail
Instance: vpls1
Learning-Domain: vlan-id 10

Interface: ge-1/3/1.10
Address: 192.0.2.2
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 83 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 830908833
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported

Interface: ge-1/3/3.10
Address: 192.0.2.3
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 97 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 2056520742
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported

Interface: ge-1/3/5.10
Address: 192.0.2.4
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 81 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 1152066227
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported

Interface: ge-1/3/7.10
Address: 192.0.2.5
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 96 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 1113200338
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported
Learning-Domain: vlan-id 20

Interface: ge-1/3/1.20
Address: 192.0.2.12
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 81 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 963205167
```

 Copyright © 2018, Juniper Networks, Inc.

```
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported

Interface: ge-1/3/3.20
Address: 192.0.2.13
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 104 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 166921538
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported

Interface: ge-1/3/5.20
Address: 192.0.2.14
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 88 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 789422835
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported

Interface: ge-1/3/7.20
Address: 192.0.2.15
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 88 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 1563649680
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported
```

show pim snooping neighbors instance

```
user@host> show pim snooping neighbors instance vpls1
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking Bit

Instance: vpls1
Learning-Domain: vlan-id 10

Interface Option Uptime Neighbor addr
ge-1/3/1.10 HPLGT 00:46:03 192.0.2.2
ge-1/3/3.10 HPLGT 00:46:03 192.0.2.3
ge-1/3/5.10 HPLGT 00:46:03 192.0.2.4
ge-1/3/7.10 HPLGT 00:46:03 192.0.2.5


Learning-Domain: vlan-id 20

Interface Option Uptime Neighbor addr
ge-1/3/1.20 HPLGT 00:46:03 192.0.2.12
ge-1/3/3.20 HPLGT 00:46:03 192.0.2.13
ge-1/3/5.20 HPLGT 00:46:03 192.0.2.14
ge-1/3/7.20 HPLGT 00:46:03 192.0.2.15
```

show pim snooping neighbors interface

```
user@host> show pim snooping neighbors interface ge-1/3/1.20
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
```

Copyright © 2018, Juniper Networks, Inc.

```
P = Hello Option DR Priority, T = Tracking Bit

Instance: vpls1
Learning-Domain: vlan-id 10
Learning-Domain: vlan-id 20

Interface Option Uptime Neighbor addr
ge-1/3/1.20 HPLGT 00:48:04 192.0.2.12
```

show pim snooping neighbors vlan-id

```
user@host> show pim snooping neighbors vlan-id 10
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking Bit

Instance: vpls1
Learning-Domain: vlan-id 10

Interface Option Uptime Neighbor addr
ge-1/3/1.10 HPLGT 00:49:12 192.0.2.2
ge-1/3/3.10 HPLGT 00:49:12 192.0.2.3
ge-1/3/5.10 HPLGT 00:49:12 192.0.2.4
ge-1/3/7.10 HPLGT 00:49:12 192.0.2.5
```


Copyright © 2018, Juniper Networks, Inc.

show pim snooping statistics

Syntax

show pim snooping statistics
<instance *instance-name*>
<interface *interface-name*>
<logical-system *logical-system-name*>
<vlan-id *vlan-id*>

Release Information

Command introduced in Junos OS Release 12.3 for MX Series 3D Universal Edge devices.
Command introduced in Junos OS Release 13.2 for M Series Multiservice Edge devices.

Description

Display Protocol Independent Multicast (PIM) snooping statistics.

Options

none—Display PIM statistics.

instance *instance-name*—(Optional) Display statistics for a specific routing instance enabled by Protocol Independent Multicast (PIM) snooping.

interface *interface-name*—(Optional) Display statistics about the specified interface for PIM snooping.

logical-system *logical-system-name*—(Optional) Display information about a particular logical system, or type 'all'.

vlan-id *vlan-identifier*—(Optional) Display PIM snooping statistics information for the specified VLAN.

Required Privilege Level

view

Related Documentation

- *PIM Snooping for VPLS*
- clear pim snooping statistics on page 1758

List of Sample Output

show pim snooping statistics on page 1940
show pim snooping statistics instance on page 1941
show pim snooping statistics interface on page 1942
show pim snooping statistics vlan-id on page 1942

Output Fields

Table 145 on page 1939 lists the output fields for the **show pim snooping statistics** command. Output fields are listed in the approximate order in which they appear.

Table 145: show pim snooping statistics Output Fields

Field Name	Field Description	Level of Output
Instance	Routing instance for PIM snooping.	All levels
Learning-Domain	Learning domain for PIM snooping.	All levels

Table 145: show pim snooping statistics Output Fields (continued)

Field Name	Field Description	Level of Output
Tx J/P messages	Total number of transmitted join/prune packets.	All levels
RX J/P messages	Total number of received join/prune packets.	All levels
Rx J/P messages -- seen	Number of join/prune packets seen but not received on the upstream interface.	All levels
Rx J/P messages -- received	Number of join/prune packets received on the downstream interface.	All levels
Rx Hello messages	Total number of received hello packets.	All levels
Rx Version Unknown	Number of packets received with an unknown version number.	All levels
Rx Neighbor Unknown	Number of packets received from an unknown neighbor.	All levels
Rx Upstream Neighbor Unknown	Number of packets received with unknown upstream neighbor information.	All levels
Rx Bad Length	Number of packets received containing incorrect length information.	All levels
Rx J/P Busy Drop	Number of join/prune packets dropped while the router is busy.	All levels
Rx J/P Group Aggregate 0	Number of join/prune packets received containing the aggregate group information.	All levels
Rx Malformed Packet	Number of malformed packets received.	All levels
Rx No PIM Interface	Number of packets received without the interface information.	All levels
Rx No Upstream Neighbor	Number of packets received without upstream neighbor information.	All levels
Rx Unknown Hello Option	Number of hello packets received with unknown options.	All levels

Sample Output

show pim snooping statistics

```
user@host> show pim snooping statistics
Instance: vpls1
Learning-Domain: vlan-id 10

Tx J/P messages 0
RX J/P messages 8
```

 Copyright © 2018, Juniper Networks, Inc.

```
Rx J/P messages -- seen 0
Rx J/P messages -- received 8
Rx Hello messages 37
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0
Rx No PIM Interface 0
Rx No Upstream Neighbor 0
Rx Bad Length 0
Rx Neighbor Unknown 0
Rx Unknown Hello Option 0
Rx Malformed Packet 0


Learning-Domain: vlan-id 20

Tx J/P messages 0
RX J/P messages 2
Rx J/P messages -- seen 0
Rx J/P messages -- received 2
Rx Hello messages 39
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0
Rx No PIM Interface 0
Rx No Upstream Neighbor 0
Rx Bad Length 0
Rx Neighbor Unknown 0
Rx Unknown Hello Option 0
Rx Malformed Packet 0
```

show pim snooping statistics instance

```
user@host> show pim snooping statistics instance vpls1
Instance: vpls1
Learning-Domain: vlan-id 10

Tx J/P messages 0
RX J/P messages 9
Rx J/P messages -- seen 0
Rx J/P messages -- received 9
Rx Hello messages 45
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0
Rx No PIM Interface 0
Rx No Upstream Neighbor 0
Rx Bad Length 0
Rx Neighbor Unknown 0
```

Copyright © 2018, Juniper Networks, Inc.

```
Rx Unknown Hello Option 0
Rx Malformed Packet 0


Learning-Domain: vlan-id 20

Tx J/P messages 0
RX J/P messages 3
Rx J/P messages -- seen 0
Rx J/P messages -- received 3
Rx Hello messages 47
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0
Rx No PIM Interface 0
Rx No Upstream Neighbor 0
Rx Bad Length 0
Rx Neighbor Unknown 0
Rx Unknown Hello Option 0
Rx Malformed Packet 0
```

show pim snooping statistics interface

```
user@host> show pim snooping statistics interface ge-1/3/1.20
Instance: vpls1
Learning-Domain: vlan-id 10
Learning-Domain: vlan-id 20

PIM Interface statistics for ge-1/3/1.20
Tx J/P messages 0
RX J/P messages 0
Rx J/P messages -- seen 0
Rx J/P messages -- received 0
Rx Hello messages 13
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0
```

show pim snooping statistics vlan-id

```
user@host> show pim snooping statistics vlan-id 10
Instance: vpls1
Learning-Domain: vlan-id 10

Tx J/P messages 0
RX J/P messages 11
Rx J/P messages -- seen 0
Rx J/P messages -- received 11
Rx Hello messages 64
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
```

 Copyright © 2018, Juniper Networks, Inc.

```
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0
Rx No PIM Interface 0
Rx No Upstream Neighbor 0
Rx Bad Length 0
Rx Neighbor Unknown 0
```

Copyright © 2018, Juniper Networks, Inc.

show pim source

List of Syntax

Syntax on page 1944
Syntax (EX Series Switch and the QFX Series) on page 1944

Syntax

show pim source
<brief | detail>
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<*source-prefix*>

Syntax (EX Series Switch and the QFX Series)

show pim source
<brief | detail>
<inet | inet6>
<instance *instance-name*>
<*source-prefix*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display information about the Protocol Independent Multicast (PIM) source reverse path forwarding (RPF) state.

Options

none—Display standard information about the PIM RPF state for all supported family addresses for all routing instances.

brief | detail—(Optional) Display the specified level of output.

inet | inet6—(Optional) Display information for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display information about the RPF state for a specific PIM-enabled routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

source-prefix—(Optional) Display the state for source RPF states in the given range.

Required Privilege Level

view

List of Sample Output

show pim source on page 1945
show pim source brief on page 1945
show pim source detail on page 1946

Copyright © 2018, Juniper Networks, Inc.

show pim source (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 1946

Output Fields Table 146 on page 1945 describes the output fields for the **show pim source** command. Output fields are listed in the approximate order in which they appear.

Table 146: show pim source Output Fields

Field Name	Field Description
Instance	Name of the routing instance.
Source	Address of the source or reverse path.
Prefix/length	Prefix and prefix length for the route used to reach the RPF address.
Upstream Protocol	Protocol toward the source address.
Upstream interface	RPF interface toward the source address. A pseudo multipoint LDP (M-LDP) interface appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.
Upstream Neighbor	Address of the RPF neighbor used to reach the source address. The multipoint LDP (M-LDP) root appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.

Sample Output

show pim source

```
user@host> show pim source
Instance: PIM.master Family: INET

Source 10.255.14.144
    Prefix 10.255.14.144/32
    Upstream interface Local
    Upstream neighbor Local

Source 10.255.70.15
    Prefix 10.255.70.15/32
    Upstream interface so-1/0/0.0
    Upstream neighbor 10.111.10.2

Instance: PIM.master Family: INET6
```

show pim source brief

The output for the **show pim source brief** command is identical to that for the **show pim source** command. For sample output, see show pim source on page 1945.

show pim source detail

```
user@host> show pim source detail
Instance: PIM.master Family: INET

Source 10.255.14.144
    Prefix 10.255.14.144/32
    Upstream interface Local
    Upstream neighbor Local
    Active groups:233.252.0.0
        233.252.0.1
        233.252.0.1

Source 10.255.70.15
    Prefix 10.255.70.15/32
    Upstream interface so-1/0/0.0
    Upstream neighbor 10.111.10.2
    Active groups:233.252.0.1

Instance: PIM.master Family: INET6
```

show pim source (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show pim source
Instance: PIM.master Family: INET

Source 10.1.1.1
    Prefix 10.1.1.1/32
    Upstream interface Local
    Upstream neighbor Local

Source 10.2.7.7
    Prefix 10.2.7.0/24
    Upstream protocol MLDP
    Upstream interface Pseudo MLDP
    Upstream neighbor MLDP LSP root <10.1.1.2>

Source 192.168.219.11
    Prefix 192.168.219.0/28
    Upstream protocol MLDP
    Upstream interface Pseudo MLDP
    Upstream neighbor via MLDP-inband
   Upstream interface fe-1/3/0.0
    Upstream neighbor 192.168.140.1
    Upstream neighbor MLDP LSP root <10.1.1.2>

Instance: PIM.master Family: INET6
Source 2001:db8::1:2:7:7
    Prefix 2001:db8::1:2:7:0/120
    Upstream protocol MLDP
    Upstream interface Pseudo MLDP
    Upstream neighbor via MLDP-inband
   Upstream interface fe-1/3/0.0
    Upstream neighbor 192.168.140.1
Upstream neighbor MLDP LSP root <10.1.1.2>
```


Copyright © 2018, Juniper Networks, Inc.

show pim statistics

List of Syntax

Syntax on page 1947
Syntax (EX Series Switch and the QFX Series) on page 1947

Syntax

show pim statistics
<inet |inet6)>
<instance *instance-name*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and the QFX Series)

show pim statistics
<inet |inet6)>
<instance *instance-name*>
<interface *interface-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Support for bidirectional PIM added in Junos OS Release 12.1.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display Protocol Independent Multicast (PIM) statistics.

Options

none—Display PIM statistics.

inet | inet6—(Optional) Display IPv4 or IPv6 PIM statistics, respectively.

instance *instance-name*—(Optional) Display statistics for a specific routing instance enabled by Protocol Independent Multicast (PIM).

interface *interface-name*—(Optional) Display statistics about the specified interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- clear pim statistics on page 1760

List of Sample Output

show pim statistics on page 1954
show pim statistics inet interface <interface-name> on page 1956
show pim statistics inet6 interface <interface-name> on page 1956
show pim statistics instance <instance-name> on page 1957
show pim statistics interface <interface-name> on page 1959

Output Fields

Table 147 on page 1948 describes the output fields for the **show pim statistics** command. Output fields are listed in the approximate order in which they appear.

Table 147: show pim statistics Output Fields

Field Name	Field Description
Instance	Name of the routing instance. This field only appears if you specify an interface, for example: • **inet interface** ***interface-name*** • **inet6 interface** ***interface-name*** • **interface** ***interface-name***
Family	Output is for IPv4 or IPv6 PIM statistics. **INET** indicates IPv4 statistics, and **INET6** indicates IPv6 statistics. This field only appears if you specify an interface, for example: • **inet interface** ***interface-name*** • **inet6 interface** ***interface-name*** • **interface** ***interface-name***
PIM statistics	PIM statistics for all interfaces or for the specified interface.
PIM message type	Message type for which statistics are displayed.
Received	Number of received statistics.
Sent	Number of messages sent of a certain type.
Rx errors	Number of received packets that contained errors.
V2 Hello	PIM version 2 hello packets.
V2 Register	PIM version 2 register packets.
V2 Register Stop	PIM version 2 register stop packets.
V2 Join Prune	PIM version 2 join and prune packets.
V2 Bootstrap	PIM version 2 bootstrap packets.
V2 Assert	PIM version 2 assert packets.
V2 Graft	PIM version 2 graft packets.
V2 Graft Ack	PIM version 2 graft acknowledgment packets.
V2 Candidate RP	PIM version 2 candidate RP packets.

Copyright © 2018, Juniper Networks, Inc.

Table 147: show pim statistics Output Fields (continued)

Field Name	Field Description
V2 State Refresh	PIM version 2 control messages related to PIM dense mode (PIM-DM) state refresh. State refresh is an extension to PIM-DM. It not supported in Junos OS.
V2 DF Election	PIM version 2 send and receive messages associated with bidirectional PIM designated forwarder election.
V1 Query	PIM version 1 query packets.
V1 Register	PIM version 1 register packets.
V1 Register Stop	PIM version 1 register stop packets.
V1 Join Prune	PIM version 1 join and prune packets.
V1 RP Reachability	PIM version 1 RP reachability packets.
V1 Assert	PIM version 1 assert packets.
V1 Graft	PIM version 1 graft packets.
V1 Graft Ack	PIM version 1 graft acknowledgment packets.
AutoRP Announce	Auto-RP announce packets.
AutoRP Mapping	Auto-RP mapping packets.
AutoRP Unknown type	Auto-RP packets with an unknown type.
Anycast Register	Auto-RP announce packets.
Anycast Register Stop	Auto-RP announce packets.
Global Statistics	Summary of PIM statistics for all interfaces.
Hello dropped on neighbor policy	Number of hello packets dropped because of a configured neighbor policy.
Unknown type	Number of PIM control packets received with an unknown type.
V1 Unknown type	Number of PIM version 1 control packets received with an unknown type.
Unknown Version	Number of PIM control packets received with an unknown version. The version is not version 1 or version 2.

Copyright © 2018, Juniper Networks, Inc.

Table 147: show pim statistics Output Fields (continued)

Field Name	Field Description
Neighbor unknown	Number of PIM control packets received (excluding PIM hello) without first receiving the hello packet.
Bad Length	Number of PIM control packets received for which the packet size does not match the PIM length field in the packet.
Bad Checksum	Number of PIM control packets received for which the calculated checksum does not match the checksum field in the packet.
Bad Receive If	Number of PIM control packets received on an interface that does not have PIM configured.
Rx Bad Data	Number of PIM control packets received that contain data for TCP Bad register packets.
Rx Intf disabled	Number of PIM control packets received on an interface that has PIM disabled.
Rx V1 Require V2	Number of PIM version 1 control packets received on an interface configured for PIM version 2.
Rx V2 Require V1	Number of PIM version 2 control packets received on an interface configured for PIM version 1.
Rx Register not RP	Number of PIM register packets received when the routing device is not the RP for the group.
Rx Register no route	Number of PIM register packets received when the RP does not have a unicast route back to the source.
Rx Register no decap if	Number of PIM register packets received when the RP does not have a de-encapsulation interface.
Null Register Timeout	Number of NULL register timeout packets.
RP Filtered Source	Number of PIM packets received when the routing device has a source address filter configured for the RP.
Rx Unknown Reg Stop	Number of register stop messages received with an unknown type.
Rx Join/Prune no state	Number of join and prune messages received for which the routing device has no state.
Rx Join/Prune on upstream if	Number of join and prune messages received on the interface used to reach the upstream routing device, toward the RP.
Rx Join/Prune for invalid group	Number of join or prune messages received for invalid multicast group addresses.

Copyright © 2018, Juniper Networks, Inc.

Table 147: show pim statistics Output Fields (continued)

Field Name	Field Description
Rx Join/Prune messages dropped	Number of join and prune messages received and dropped.
Rx sparse join for dense group	Number of PIM sparse mode join messages received for a group that is configured for dense mode.
Rx Graft/Graft Ack no state	Number of graft and graft acknowledgment messages received for which the router or switch has no state.
Rx Graft on upstream if	Number of graft messages received on the interface used to reach the upstream routing device, toward the RP.
Rx CRP not BSR	Number of BSR messages received in which the PIM message type is Candidate-RP-Advertisement, not Bootstrap.
Rx BSR when BSR	Number of BSR messages received in which the PIM message type is Bootstrap.
Rx BSR not RPF if	Number of BSR messages received on an interface that is not the RPF interface.
Rx unknown hello opt	Number of PIM hello packets received with options that Junos OS does not support.
Rx data no state	Number of PIM control packets received for which the routing device has no state for the data type.
Rx RP no state	Number of PIM control packets received for which the routing device has no state for the RP.
Rx aggregate	Number of PIM aggregate MDT packets received.
Rx malformed packet	Number of PIM control packets received with a malformed IP unicast or multicast address family.
No RP	Number of PIM control packets received with no RP address.
No register encap if	Number of PIM register packets received when the first-hop routing device does not have an encapsulation interface.
No route upstream	Number of PIM control packets received when the routing device does not have a unicast route to the the interface used to reach the upstream routing device, toward the RP.
Nexthop Unusable	Number of PIM control packets with an unusable nexthop. A path can be unusable if the route is hidden or the link is down.
RP mismatch	Number of PIM control packets received for which the routing device has an RP mismatch.

Copyright © 2018, Juniper Networks, Inc.

Table 147: show pim statistics Output Fields (continued)

Field Name	Field Description
RP mode mismatch	RP mode (sparse or bidirectional) mismatches encountered when processing join and prune messages.
RPF neighbor unknown	Number of PIM control packets received for which the routing device has an unknown RPF neighbor for the source.
Rx Joins/Prunes filtered	The number of join and prune messages filtered because of configured route filters and source address filters.
Tx Joins/Prunes filtered	The number of join and prune messages filtered because of configured route filters and source address filters.
Embedded-RP invalid addr	Number of packets received with an invalid embedded RP address in PIM join messages and other types of messages sent between routing domains.
Embedded-RP limit exceed	Number of times the limit configured with the **maximum-rps** statement is exceeded. The **maximum-rps** statement limits the number of embedded RPs created in a specific routing instance. The range is from 1 through 500. The default is 100.
Embedded-RP added	Number of packets in which the embedded RP for IPv6 is added. The following receive events trigger extraction of an IPv6 embedded RP address on the routing device: • Multicast Listener Discovery (MLD) report for an embedded RP multicast group address • PIM join message with an embedded RP multicast group address • Static embedded RP multicast group address associated with an interface • Packets sent to an embedded RP multicast group address received on the DR An embedded RP node discovered through these receive events is added if it does not already exist on the routing platform.
Embedded-RP removed	Number of packets in which the embedded RP for IPv6 is removed. The embedded RP is removed whenever all PIM join states using this RP are removed or the configuration changes to remove the embedded RP feature.
Rx Register msgs filtering drop	Number of received register messages dropped because of a filter configured for PIM register messages.
Tx Register msgs filtering drop	Number of register messages dropped because of a filter configured for PIM register messages.
Rx Bidir Join/Prune on non-Bidir if	Error counter for join and prune messages received on non-bidirectional PIM interfaces.

 Copyright © 2018, Juniper Networks, Inc.

Table 147: show pim statistics Output Fields (continued)

Field Name	Field Description
Rx Bidir Join/Prune on non-DF if	Error counter for join and prune messages received on non-designated forwarder interfaces.
V4 (S,G) Maximum	Maximum number of (S,G) IPv4 multicast routes accepted for the VPN routing and forwarding (VRF) routing instance. If this number is met, additional (S,G) entries are not accepted.
V4 (S,G) Accepted	Number of accepted (S,G) IPv4 multicast routes.
V4 (S,G) Threshold	Threshold at which a warning message is logged (percentage of the maximum number of (S,G) IPv4 multicast routes accepted by the device).
V4 (S,G) Log Interval	Time (in seconds) between consecutive log messages.
V6 (S,G) Maximum	Maximum number of (S,G) IPv6 multicast routes accepted for the VPN routing and forwarding (VRF) routing instance. If this number is met, additional (S,G) entries are not accepted.
V6 (S,G) Accepted	Number of accepted (S,G) IPv6 multicast routes.
V6 (S,G) Threshold	Threshold at which a warning message is logged (percentage of the maximum number of (S,G) IPv6 multicast routes accepted by the device).
V6 (S,G) Log Interval	Time (in seconds) between consecutive log messages.
V4 (grp-prefix, RP) Maximum	Maximum number of group-to-rendezvous point (RP) IPv4 multicast mappings accepted for the VRF routing instance. If this number is met, additional mappings are not accepted.
V4 (grp-prefix, RP) Accepted	Number of accepted group-to-RP IPv4 multicast mappings.
V4 (grp-prefix, RP) Threshold	Threshold at which a warning message is logged (percentage of the maximum number of group-to-RP IPv4 multicast mappings accepted by the device).
V4 (grp-prefix, RP) Log Interval	Time (in seconds) between consecutive log messages.
V6 (grp-prefix, RP) Maximum	Maximum number of group-to RP IPv6 multicast mappings accepted for the VRF routing instance. If this number is met, additional mappings are not accepted.
V6 (grp-prefix, RP) Accepted	Number of accepted group-to-RP IPv6 multicast mappings.

Copyright © 2018, Juniper Networks, Inc.

Table 147: show pim statistics Output Fields (continued)

Field Name	Field Description
V6 (grp-prefix, RP) Threshold	Threshold at which a warning message is logged (percentage of the maximum number of group-to-RP IPv6 multicast mappings accepted by the device).
V6 (grp-prefix, RP) Log Interval	Time (in seconds) between consecutive log messages.
V4 Register Maximum	Maximum number of IPv4 PIM registers accepted for the VRF routing instance. If this number is met, additional PIM registers are not accepted. You configure the register limits on the RP.
V4 Register Accepted	Number of accepted IPv4 PIM registers.
V4 Register Threshold	Threshold at which a warning message is logged (percentage of the maximum number of IPv4 PIM registers accepted by the device).
V4 Register Log Interval	Time (in seconds) between consecutive log messages.
V6 Register Maximum	Maximum number of IPv6 PIM registers accepted for the VRF routing instance. If this number is met, additional PIM registers are not accepted. You configure the register limits on the RP.
V6 Register Accepted	Number of accepted IPv6 PIM registers.
V6 Register Threshold	Threshold at which a warning message is logged (percentage of the maximum number of IPv6 PIM registers accepted by the device).
V6 Register Log Interval	Time (in seconds) between consecutive log messages.
(*,G) Join drop due to SSM range check	PIM join messages that are dropped because the multicast addresses are outside of the SSM address range of 232.0.0.0 through 232.255.255.255. You can extend the accepted SSM address range by configuring the **ssm-groups** statement.

Sample Output

show pim statistics

```
user@host> show pim statistics
PIM Message type        Received       Sent  Rx errors
V2 Hello                      15         32          0
V2 Register                    0        362          0
V2 Register Stop             483          0          0
V2 Join Prune                 18        518          0
V2 Bootstrap                   0          0          0
V2 Assert                      0          0          0
V2 Graft                       0          0          0
```

 Copyright © 2018, Juniper Networks, Inc.

```
V2 Graft Ack                0          0          0
V2 Candidate RP             0          0          0
V2 State Refresh            0          0          0
V2 DF Election              0          0          0
V1 Query                    0          0          0
V1 Register                 0          0          0
V1 Register Stop            0          0          0
V1 Join Prune               0          0          0
V1 RP Reachability          0          0          0
V1 Assert                   0          0          0
V1 Graft                    0          0          0
V1 Graft Ack                0          0          0
AutoRP Announce             0          0          0
AutoRP Mapping              0          0          0
AutoRP Unknown type         0
Anycast Register            0          0          0
Anycast Register Stop       0          0          0

Global Statistics

Hello dropped on neighbor policy              0
Unknown type                                  0
V1 Unknown type                               0
Unknown Version                               0
ipv4 BSR pkt drop due to excessive rate       0
ipv6 BSR pkt drop due to excessive rate       0
Neighbor unknown                              0
Bad Length                                    0
Bad Checksum                                  0
Bad Receive If                                0
Rx Bad Data                                   0
Rx Intf disabled                              0
Rx V1 Require V2                              0
Rx V2 Require V1                              0
Rx Register not RP                            0
Rx Register no route                          0
Rx Register no decap if                       0
Null Register Timeout                         0
RP Filtered Source                            0
Rx Unknown Reg Stop                           0
Rx Join/Prune no state                        0
Rx Join/Prune on upstream if                  0
Rx Join/Prune for invalid group               5
Rx Join/Prune messages dropped                0
Rx sparse join for dense group                0
Rx Graft/Graft Ack no state                   0
Rx Graft on upstream if                       0
Rx CRP not BSR                                0
Rx BSR when BSR                               0
Rx BSR not RPF if                             0
Rx unknown hello opt                          0
Rx data no state                              0
Rx RP no state                                0
Rx aggregate                                  0
Rx malformed packet                           0
Rx illegal TTL                                0
Rx illegal destination address                0
No RP                                         0
No register encap if                          0
No route upstream                             0
Nexthop Unusable                              0
```

```
RP mismatch                               0
RP mode mismatch                          0
RPF neighbor unknown                      0
Rx Joins/Prunes filtered                  0
Tx Joins/Prunes filtered                  0
Embedded-RP invalid addr                  0
Embedded-RP limit exceed                  0
Embedded-RP added                         0
Embedded-RP removed                       0
Rx Register msgs filtering drop           0
Tx Register msgs filtering drop           0
Rx Bidir Join/Prune on non-Bidir if       0
Rx Bidir Join/Prune on non-DF if          0
(*,G) Join drop due to SSM range check    0
```

Sample Output

show pim statistics inet interface <interface-name>

```
user@host> show pim statistics inet interface ge-0/3/0.0
Instance: PIM.master Family: INET

PIM Interface statistics for ge-0/3/0.0

PIM Message type        Received       Sent  Rx errors
V2 Hello                       0          4          0
V2 Register                    0          0          0
V2 Register Stop               0          0          0
V2 Join Prune                  0          0          0
V2 Bootstrap                   0          0          0
V2 Assert                      0          0          0
V2 Graft                       0          0          0
V2 Graft Ack                   0          0          0
V2 Candidate RP                0          0          0
V1 Query                       0          0          0
V1 Register                    0          0          0
V1 Register Stop               0          0          0
V1 Join Prune                  0          0          0
V1 RP Reachability             0          0          0
V1 Assert                      0          0          0
V1 Graft                       0          0          0
V1 Graft Ack                   0          0          0
AutoRP Announce                0          0          0
AutoRP Mapping                 0          0          0
AutoRP Unknown type            0
Anycast Register               0          0          0
Anycast Register Stop          0          0          0
```

Sample Output

show pim statistics inet6 interface <interface-name>

```
user@host> show pim statistics inet6 interface ge-0/3/0.0
Instance: PIM.master Family: INET6

PIM Interface statistics for ge-0/3/0.0

PIM Message type        Received       Sent  Rx errors
V2 Hello                       0          4          0
V2 Register                    0          0          0
```


Copyright © 2018, Juniper Networks, Inc.

```
V2 Register Stop               0          0          0
V2 Join Prune                  0          0          0
V2 Bootstrap                   0          0          0
V2 Assert                      0          0          0
V2 Graft                       0          0          0
V2 Graft Ack                   0          0          0
V2 Candidate RP                0          0          0
Anycast Register               0          0          0
Anycast Register Stop          0          0          0
```

show pim statistics instance <instance-name>

```
user@host> show pim statistics instance VPN-A
PIM Message type        Received       Sent  Rx errors
V2 Hello                      31         37          0
V2 Register                    0          0          0
V2 Register Stop               0          0          0
V2 Join Prune                  0         16          0
V2 Bootstrap                   0          0          0
V2 Assert                      0          0          0
V2 Graft                       0          0          0
V2 Graft Ack                   0          0          0
V2 Candidate RP                0          0          0
V2 State Refresh               0          0          0
V2 DF Election                 0          0          0
V1 Query                       0          0          0
V1 Register                    0          0          0
V1 Register Stop               0          0          0
V1 Join Prune                  0          0          0
V1 RP Reachability             0          0          0
V1 Assert                      0          0          0
V1 Graft                       0          0          0
V1 Graft Ack                   0          0          0
AutoRP Announce                0          0          0
AutoRP Mapping                 0          0          0
AutoRP Unknown type            0
Anycast Register               0          0          0
Anycast Register Stop          0          0          0

Global Statistics

Hello dropped on neighbor policy               0
Unknown type                                   0
V1 Unknown type                                0
Unknown Version                                0
Neighbor unknown                               0
Bad Length                                     0
Bad Checksum                                   0
Bad Receive If                                 0
Rx Bad Data                                    0
Rx Intf disabled                               0
Rx V1 Require V2                               0
Rx V2 Require V1                               0
Rx Register not RP                             0
Rx Register no route                           0
Rx Register no decap if                        0
Null Register Timeout                          0
RP Filtered Source                             0
Rx Unknown Reg Stop                            0
Rx Join/Prune no state                         0
Rx Join/Prune on upstream if                   0
```

```
Rx Join/Prune for invalid group             0
Rx Join/Prune messages dropped              0
Rx sparse join for dense group              0
Rx Graft/Graft Ack no state                 0
Rx Graft on upstream if                     0
Rx CRP not BSR                              0
Rx BSR when BSR                             0
Rx BSR not RPF if                           0
Rx unknown hello opt                        0
Rx data no state                            0
Rx RP no state                              0
Rx aggregate                                0
Rx malformed packet                         0
Rx illegal TTL                              0
Rx illegal destination address              0
No RP                                       0
No register encap if                        0
No route upstream                          28
Nexthop Unusable                            0
RP mismatch                                 0
RP mode mismatch                            0
RPF neighbor unknown                        0
Rx Joins/Prunes filtered                    0
Tx Joins/Prunes filtered                    0
Embedded-RP invalid addr                    0
Embedded-RP limit exceed                    0
Embedded-RP added                           0
Embedded-RP removed                         0
Rx Register msgs filtering drop             0
Tx Register msgs filtering drop             0
Rx Bidir Join/Prune on non-Bidir if         0
Rx Bidir Join/Prune on non-DF if            0
V4 (S,G) Maximum                           10
V4 (S,G) Accepted                           9
V4 (S,G) Threshold                         80
V4 (S,G) Log Interval                      80
V6 (S,G) Maximum                            8
V6 (S,G) Accepted                           8
V6 (S,G) Threshold                         50
V6 (S,G) Log Interval                     100
V4 (grp-prefix, RP) Maximum               100
V4 (grp-prefix, RP) Accepted                5
V4 (grp-prefix, RP) Threshold              80
V4 (grp-prefix, RP) Log Interval           10
V6 (grp-prefix, RP) Maximum                20
V6 (grp-prefix, RP) Accepted                0
V6 (grp-prefix, RP) Threshold              90
V6 (grp-prefix, RP) Log Interval           20
V4 Register Maximum                       100
V4 Register Accepted                       10
V4 Register Threshold                      80
V4 Register Log Interval                   10
V6 Register Maximum                        20
V6 Register Accepted                        0
V6 Register Threshold                      90
V6 Register Log Interval                   20
(*,G) Join drop due to SSM range check      0
```

 Copyright © 2018, Juniper Networks, Inc.

Sample Output

show pim statistics interface <interface-name>

```
user@host> show pim statistics interface ge-0/3/0.0
Instance: PIM.master Family: INET

PIM Interface statistics for ge-0/3/0.0

PIM Message type        Received       Sent  Rx errors
V2 Hello                       0          3          0
V2 Register                    0          0          0
V2 Register Stop               0          0          0
V2 Join Prune                  0          0          0
V2 Bootstrap                   0          0          0
V2 Assert                      0          0          0
V2 Graft                       0          0          0
V2 Graft Ack                   0          0          0
V2 Candidate RP                0          0          0
V1 Query                       0          0          0
V1 Register                    0          0          0
V1 Register Stop               0          0          0
V1 Join Prune                  0          0          0
V1 RP Reachability             0          0          0
V1 Assert                      0          0          0
V1 Graft                       0          0          0
V1 Graft Ack                   0          0          0
AutoRP Announce                0          0          0
AutoRP Mapping                 0          0          0
AutoRP Unknown type            0
Anycast Register               0          0          0
Anycast Register Stop          0          0          0

Instance: PIM.master Family: INET6

PIM Interface statistics for ge-0/3/0.0

PIM Message type        Received       Sent  Rx errors
V2 Hello                       0          3          0
V2 Register                    0          0          0
V2 Register Stop               0          0          0
V2 Join Prune                  0          0          0
V2 Bootstrap                   0          0          0
V2 Assert                      0          0          0
V2 Graft                       0          0          0
V2 Graft Ack                   0          0          0
V2 Candidate RP                0          0          0
Anycast Register               0          0          0
Anycast Register Stop          0          0          0
```

show sap listen

Syntax

show sap listen
<brief | detail>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display the addresses that the router is listening to in order to receive multicast Session Announcement Protocol (SAP) session announcements.

Options

none—Display standard information about the addresses that the router is listening to in order to receive multicast SAP session announcements.

brief | detail—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show sap listen on page 1960
show sap listen brief on page 1961
show sap listen detail on page 1961

Output Fields

Table 148 on page 1960 describes the output fields for the **show sap listen** command. Output fields are listed in the approximate order in which they appear.

Table 148: show sap listen Output Fields

Field Name	Field Description
Group address	Address of the group that the local router is listening to for SAP messages.
Port	UDP port number used for SAP.

Sample Output

show sap listen

```
user@host> show sap listen
Group address   Port
224.2.127.254   9875
239.255.255.255 9875
```


Copyright © 2018, Juniper Networks, Inc.

show sap listen brief

The output for the **show sap listen brief** command is identical to that for the **show sap listen** command. For sample output, see show sap listen on page 1960.

show sap listen detail

The output for the **show sap listen detail** command is identical to that for the **show sap listen** command. For sample output, see show sap listen on page 1960.

Copyright © 2018, Juniper Networks, Inc.

test msdp

Syntax

test msdp (dependent-peers *prefix* | rpf-peer *originator*)
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Find Multicast Source Discovery Protocol (MSDP) peers.

Options

dependent-peers *prefix*—Find downstream dependent MSDP peers.

rpf-peer *originator*—Find the MSDP reverse-path-forwarding (RPF) peer for the originator.

instance *instance-name*—(Optional) Find MDSP peers for the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

test msdp dependent-peers on page 1962

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

test msdp dependent-peers

```
user@host> test msdp dependent-peers 10.0.0.1/24
```


Copyright © 2018, Juniper Networks, Inc.

CHAPTER 20

IPv6 Operational Commands

clear ipv6 neighbors

Syntax

clear ipv6 neighbors
<all | host *hostname*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.3 for EX Series switches.
Command introduced in Junos OS Release 12.2 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Clear IPv6 neighbor cache information.

Options

none—Clear all IPv6 neighbor cache information.

all—(Optional) Clear all IPv6 neighbor cache information.

host *hostname*—(Optional) Clear the information for the specified IPv6 neighbors.

Required Privilege Level

view

Related Documentation

- show ipv6 neighbors on page 1966

List of Sample Output

clear ipv6 neighbors on page 1964

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ipv6 neighbors

```
user@host> clear ipv6 neighbors
```


Copyright © 2018, Juniper Networks, Inc.

clear ipv6 router-advertisement

Syntax

clear ipv6 router-advertisement
<interface *interface*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced before Junos OS Release 7.4.

Description

Clear IPv6 router advertisement counters.

Options

none—Clear IPv6 router advertisement counters for all interfaces.

interface *interface*—(Optional) Clear IPv6 router advertisement counters for the specified interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- show ipv6 router-advertisement on page 1968

List of Sample Output

clear ipv6 router-advertisement on page 1965

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ipv6 router-advertisement

```
user@host> clear ipv6 router-advertisement
```

Copyright © 2018, Juniper Networks, Inc.

show ipv6 neighbors

Syntax show ipv6 neighbors

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.3 for EX Series switches.
Command introduced in Junos OS Release 12.2 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Display information about the IPv6 neighbor cache.

NOTE: Starting with Junos OS Release 16.1, **show ipv6 neighbors** command does not display the underlying ifl information if **enhanced-convergence** statement at **[edit irb unit *unit-number*]** hierarchy level and **enhanced-ip** statement at **[edit chassis network-services]** hierarchy level is configured for the destination interface IRB.

Options This command has no options.

Required Privilege Level view

Related Documentation

- clear ipv6 neighbors on page 1964

List of Sample Output show ipv6 neighbors on page 1967
show ipv6 neighbors on page 1967

Output Fields Table 149 on page 1966 describes the output fields for the **show ipv6 neighbors** command. Output fields are listed in the approximate order in which they appear.

Table 149: show ipv6 neighbors Output Fields

Field Name	Field Description
IPv6 Address	Name of the IPv6 interface.
Linklayer Address	Link-layer address.
State	State of the link: **up**, **down**, **incomplete**, **reachable**, **stale**, or **unreachable**.
Exp	Number of seconds until the entry expires.
Rtr	Whether the neighbor is a routing device: **yes** or **no**.

 Copyright © 2018, Juniper Networks, Inc.

Table 149: show ipv6 neighbors Output Fields (continued)

Field Name	Field Description
Secure	Whether this entry was created using the Secure Neighbor Discovery (SEND) protocol: **yes** or **no**.
Interface	Name of the interface.

Sample Output

show ipv6 neighbors

```
user@host> show ipv6 neighbors
IPv6 Address                 Linklayer Address  State       Exp Rtr Secure
Interface
2001:db8:0:1:2a0:a514:0:24c  00:05:85:8f:c8:bd  stale       546 yes no
fe-1/2/0.1
fe80::2a0:a514:0:24c         00:05:85:8f:c8:bd  stale       258 yes no
fe-1/2/0.1
fe80::2a0:a514:0:64c         00:05:85:8f:c8:bd  stale       111 yes no
fe-1/2/1.5
fe80::2a0:a514:0:a4c         00:05:85:8f:c8:bd  stale       327 yes no
fe-1/2/2.9
```

show ipv6 neighbors

The command displaying the underlying l2 ifl information when **enhanced-convergence** statement and **enhanced-ip** statement is not configured.

```
IPv6 Address                 Linklayer Address  State       Exp Rtr Secure
Interface
23::23:0:0:2                00:00:23:00:00:02  reachable   0   no  no      irb.0
 [xe-2/2/0.0]
```

The command not displaying the underlying l2 ifl information when **enhanced-convergence** statement and **enhanced-ip** statement is configured.

```
IPv6 Address                 Linklayer Address  State       Exp Rtr Secure
Interface
23::23:0:0:2                00:00:23:00:00:02  reachable   0   no  no      irb.0
```

Copyright © 2018, Juniper Networks, Inc.

show ipv6 router-advertisement

Syntax

show ipv6 router-advertisement
<conflicts>
<interface *interface*>
<logical-system (all | *logical-system-name*)>
<prefix *prefix/prefix length*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.2 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display information about IPv6 router advertisements, including statistics about messages sent and received on interfaces, and information received from advertisements from other routers.

Options

none—Display all IPv6 router advertisement information for all interfaces.

conflicts—(Optional) Display only the IPv6 router advertisement information that is conflicting.

interface *interface*—(Optional) Display IPv6 router advertisement information for the specified interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

prefix *prefix/prefix length*—(Optional) Display IPv6 router advertisement information for the specified prefix.

Additional Information

The display identifies conflicting information by enclosing the value the router is advertising in brackets.

Required Privilege Level

view

Related Documentation

- clear ipv6 router-advertisement on page 1965

List of Sample Output

show ipv6 router-advertisement on page 1969
show ipv6 router-advertisement conflicts on page 1970
show ipv6 router-advertisement prefix on page 1970

Output Fields

Table 150 on page 1969 describes the output fields for the **show ipv6 router-advertisement** command. Output fields are listed in the approximate order in which they appear.

 Copyright © 2018, Juniper Networks, Inc.

Table 150: show ipv6 router-advertisement Output Fields

Field Name	Field Description
Interface	Name of the interface.
Advertisements sent	Number of router advertisements sent and the elapsed time since they were sent.
Solicits received	Number of solicitation messages received.
Advertisements received	Number of router advertisements received.
Advertisements from	Names of interfaces from which router advertisements have been received and the elapsed time since the last one was received.
Managed	Managed address configuration flag: **0** (stateless) or **1** (stateful).
Other configuration	Other stateful configuration flag: **0** (stateless) or **1** (stateful).
Reachable time	Time that a node identifies a neighbor as reachable after receiving a reachability confirmation, in milliseconds.
Default lifetime	Default lifetime, in seconds: from 0 seconds to 18.2 hours. A setting of 0 indicates that the router is not a default router.
Retransmit timer	Time between retransmitted Neighbor Solicitation messages, in milliseconds.
Current hop limit	Configured current hop limit.
Prefix	Name and length of the prefix.
Valid lifetime	How long the prefix remains valid for onlink determination.
Preferred lifetime	How long the prefix generated by stateless autoconfiguration remains preferred.
On link	Onlink flag: **0** (not onlink) or **1** (onlink).
Autonomous	Autonomous address configuration flag: **0** (not autonomous) or **1** (autonomous).

Sample Output

show ipv6 router-advertisement

```
user@host> show ipv6 router-advertisement
Interface: fe-0/1/1.0
  Advertisements sent: 0
  Solicits received: 0
  Advertisements received: 0
Interface: fxp0.0
  Advertisements sent: 0
  Solicits received: 0
  Advertisements received: 1
```

Copyright © 2018, Juniper Networks, Inc.

```
  Advertisement from fe80::2d0:b7ff:fe1e:7b0e, heard 00:00:13 ago
    Managed: 0
    Other configuration: 0 [1]
      Reachable time: 0 ms
      Default lifetime: 1800 sec
      Retransmit timer: 0 ms
      Current hop limit: 64
```

show ipv6 router-advertisement conflicts

```
user@host> show ipv6 router-advertisement conflicts
Interface: fxp0.0
  Advertisement from fe80::2d0:b7ff:fe1e:7b0e, heard 00:01:08 ago
    Other configuration: 0 [1]
```

show ipv6 router-advertisement prefix

```
user@host> show ipv6 router-advertisement prefix 2001:db8:8040::/16
Interface: fe-0/1/3.0
  Advertisements sent: 3, last sent 00:04:11 ago
  Solicits received: 0
  Advertisements received: 3
  Advertisement from fe80::290:69ff:fe9a:5403, heard 00:00:05 ago
    Managed: 0
    Other configuration: 0
    Reachable time: 0 ms
    Default lifetime: 180 sec [1800 sec]
    Retransmit timer: 0 ms
    Current hop limit: 64
    Prefix: 2001:db8:8040:1::/64
      Valid lifetime: 2592000 sec
      Preferred lifetime: 604800 sec
      On link: 1
      Autonomous: 1
```

 Copyright © 2018, Juniper Networks, Inc.

CHAPTER 21

IS-IS Operational Commands

clear isis adjacency

List of Syntax

Syntax on page 1972
Syntax (EX Series Switches and QFX Series) on page 1972

Syntax

clear isis adjacency
<all>
<instance *instance-name*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>
<*neighbor*>

Syntax (EX Series Switches and QFX Series)

clear isis adjacency
<all>
<instance *instance-name*>
<interface *interface-name*>
<*neighbor*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 12.1 for the QFX Series.
all option introduced in Junos OS Release 14.2.

Description

Remove entries from the IS-IS adjacency database.

Options

all—Remove all entries from the adjacency database.

instance *instance-name*—(Optional) Clear all adjacencies for the specified routing instance only.

interface *interface-name*—(Optional) Clear all adjacencies for the specified interface only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

neighbor—(Optional) Clear adjacencies for the specified neighbor only.

Required Privilege Level

clear

Related Documentation

- show isis adjacency on page 1980

List of Sample Output

clear isis adjacency on page 1973
clear isis adjacency all on page 1973

Output Fields

See show isis adjacency for an explanation of output fields.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

clear isis adjacency

The following sample output displays IS-IS adjacency database information before and after the **clear isis adjacency** command is entered:

```
user@host> show isis adjacency
IS-IS adjacency database:
Interface     System         L State        Hold (secs) SNPA
so-1/0/0.0    karakul        3 Up                    26
so-1/1/3.0    1921.6800.5080 3 Up                    23
so-5/0/0.0    1921.6800.5080 3 Up                    19
```

```
user@host> clear isis adjacency karakul
```

```
user@host> show isis adjacency
IS-IS adjacency database:
Interface     System         L State        Hold (secs) SNPA
so-1/0/0.0    karakul        3 Initializing          26
so-1/1/3.0    1921.6800.5080 3 Up                    24
so-5/0/0.0    1921.6800.5080 3 Up                    21
```

clear isis adjacency all

```
user@host> clear isis adjacency all
IS-IS adjacency database:
Interface     System         L State        Hold (secs) SNPA
so-1/0/0.0    karakul        3 Initializing          26
so-1/1/3.0    1921.6800.5080 3 Initializing          24
so-5/0/0.0    1921.6800.5080 3 Initializing          21
```

clear isis database

List of Syntax

Syntax on page 1974
Syntax (EX Series Switches and QFX Series) on page 1974

Syntax

clear isis database
<all>
<*entries*>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches and QFX Series)

clear isis database
<all>
<*entries*>
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Command introduced in 15.1X53-D30 for QFX10002 switch.

Description

Remove the entries from the IS-IS link-state database, which contains prefixes and topology information.

Options

all—Remove all entries from the IS-IS link-state database for all routing instances.

entries—(Optional) Name of the database entry.

instance *instance-name*—(Optional) Clear all entries for the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Related Documentation

- show isis database on page 1997

List of Sample Output

clear isis database on page 1975

Output Fields

See show isis database for an explanation of output fields.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

clear isis database

The following sample output displays IS-IS link-state database information before and after the **clear isis database all** command is entered:

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                  Sequence Checksum Lifetime (secs)
crater.00-00                0x12   0x84dd           1139
  1 LSPs
IS-IS level 2 link-state database:
LSP ID                  Sequence Checksum Lifetime (secs)
crater.00-00                0x19   0xe92c           1134
badlands.00-00              0x16   0x1454            985
carlsbad.00-00              0x33   0x220b           1015
ranier.00-00                0x2e   0xfc31           1007
1921.6800.5066.00-00        0x11   0x7313            566
1921.6800.5067.00-00        0x14   0xd9d4            939
  6 LSPs
```

```
user@host> clear isis database all
```

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                  Sequence Checksum Lifetime (secs)

IS-IS level 2 link-state database:
LSP ID                  Sequence Checksum Lifetime (secs)
```

Copyright © 2018, Juniper Networks, Inc.

clear isis overload

List of Syntax
Syntax on page 1976
Syntax (EX Series Switches and QFX Series) on page 1976

Syntax
clear isis overload
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches and QFX Series)
clear isis overload
<instance *instance-name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description
Reset the IS-IS dynamic overload bit. This command can appear to not work, continuing to display **overload** after execution. The bit is reset only if the root cause is corrected by configuration remotely or locally.

When other routers detect that the overload bit is set, they do not use this routing device for transit traffic, but they do use it for packets destined to the overloaded routing device's directly connected networks and IP prefixes.

Options
none—Reset the IS-IS dynamic overload bit.

instance *instance-name*—(Optional) Reset the IS-IS dynamic overload bit for the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
clear

Related Documentation
- show isis database on page 1997

List of Sample Output
clear isis overload on page 1977

Output Fields
See show isis database for an explanation of output fields.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

clear isis overload

The following sample output displays IS-IS database information before and after the **clear isis overload** command is entered:

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                      Sequence Checksum Lifetime Attributes
pro3-c.00-00                     0x4   0x10db     1185 L1 L2 Overload

  1 LSPs
IS-IS level 2 link-state database:
LSP ID                      Sequence Checksum Lifetime Attributes
pro3-c.00-00                     0x5   0x429f     1185 L1 L2 Overload


pro2-a.00-00                   0x91e   0x2589      874 L1 L2
pro2-a.02-00                     0x1    0xcbc      874 L1 L2
  3 LSPs

user@host> clear isis overload


user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                      Sequence Checksum Lifetime Attributes
pro3-c.00-00                     0xa   0x429e     1183 L1 L2
  1 LSPs

IS-IS level 2 link-state database:
LSP ID                      Sequence Checksum Lifetime Attributes
pro3-c.00-00                     0xc   0x9c39     1183 L1 L2
pro2-a.00-00                   0x91e   0x2589      783 L1 L2
pro2-a.02-00                     0x1    0xcbc      783 L1 L2
  3 LSPs
```

clear isis statistics

List of Syntax

Syntax on page 1978
Syntax (EX Series Switches and QFX Series) on page 1978

Syntax

clear isis statistics
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches and QFX Series)

clear isis statistics
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Set statistics about IS-IS traffic to zero.

Options

none—Set IS-IS traffic statistics to zero for all routing instances.

instance *instance-name*—(Optional) Set IS-IS traffic statistics to zero for the specified routing instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- show isis statistics on page 2032

List of Sample Output

clear isis statistics on page 1978

Output Fields

See show isis statistics for an explanation of output fields.

Sample Output

clear isis statistics

The following sample output displays IS-IS statistics before and after the **clear isis statistics** command is entered:

```
user@host> show isis statistics
IS-IS statistics for merino:

PDU type     Received  Processed      Drops       Sent      Rexmit
```


Copyright © 2018, Juniper Networks, Inc.

```
LSP              12793      12793          0      8666        719
IIH             116751     116751          0    118834          0
CSNP            203956     203956          0    204080          0
PSNP              7356       7350          6      8635          0
Unknown              0          0          0         0          0
Totals          340856     340850          6    340215        719

Total packets received: 340856 Sent: 340934

SNP queue length:          0 Drops:          0
LSP queue length:          0 Drops:          0

SPF runs:               1064
Fragments rebuilt:      1087
LSP regenerations:       436
Purges initiated:          0
```

user@host> **clear isis statistics**

user@host> **show isis statistics**

```
IS-IS statistics for merino:

PDU type     Received  Processed      Drops       Sent     Rexmit
LSP                 0          0          0          0          0
IIH                 3          3          0          3          0
CSNP                2          2          0          4          0
PSNP                0          0          0          0          0
Unknown             0          0          0          0          0
Totals              5          5          0          7          0

Total packets received: 5 Sent: 7

SNP queue length:          0 Drops:          0
LSP queue length:          0 Drops:          0

SPF runs:                  0
Fragments rebuilt:         0
LSP regenerations:         0
Purges initiated:          0
```

show isis adjacency

List of Syntax

Syntax on page 1980
Syntax (EX Series Switches and QFX Series) on page 1980

Syntax

show isis adjacency
<system-id>
<brief | detail | extensive>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches and QFX Series)

show isis adjacency
<system-id>
<brief | detail | extensive>
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display information about IS-IS neighbors.

Options

none—Display standard information about IS-IS neighbors for all routing instances.

system id—(Optional) Display information about IS-IS neighbors for the specified intermediate system.

brief | detail | extensive—(Optional) Display standard information about IS-IS neighbors with the specified level of output.

instance *instance-name*—(Optional) Display information about IS-IS neighbors for the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Display information about IS-IS neighbors for all logical systems or for a particular logical system.

Required Privilege Level

view

Related Documentation

• clear isis adjacency on page 1972

List of Sample Output

show isis adjacency on page 1983
show isis adjacency brief on page 1983
show isis adjacency detail on page 1984
show isis adjacency extensive on page 1984

Copyright © 2018, Juniper Networks, Inc.

Output Fields Table 151 on page 1981 describes the output fields for the **show isis adjacency** command. Output fields are listed in the approximate order in which they appear.

Table 151: show isis adjacency Output Fields

Field Name	Field Description	Level of Output
Interface	Interface through which the neighbor is reachable.	All levels
System	System identifier (**sysid**), displayed as a name, if possible.	**brief**
L or Level	Level: • **1**—Level 1 only • **2**—Level 2 only • **3**—Level 1 and Level 2 An exclamation point (**!**) preceding the level number indicates that the adjacency is missing an IP address.	All levels
State	State of the adjacency: **Up**, **Down**, **New**, **One-way**, **Initializing**, or **Rejected**.	All levels
Hold (secs)	Remaining hold time of the adjacency.	**brief**
SNPA	Subnetwork point of attachment (MAC address of the next hop).	**brief**
Expires in	How long until the adjacency expires, in seconds.	**detail**
Priority	Priority to become the designated intermediate system.	**detail extensive**
Up/Down transitions	Count of adjacency status changes from **Up** to **Down** or from **Down** to **Up**.	**detail**
Last transition	Time of the last **Up/Down** transition.	**detail**
Circuit type	Bit mask of levels on this interface: **1**=Level 1 router; **2**=Level 2 router; **3**=both Level 1 and Level 2 router.	**detail**
Speaks	Protocols supported by this neighbor.	**detail extensive**
MAC address	MAC address of the interface.	**detail extensive**
Topologies	Supported topologies.	**detail extensive**
Restart capable	Whether a neighbor is capable of graceful restart: **Yes** or **No**.	**detail extensive**
Adjacency advertisement: Advertise	This routing device has signaled to advertise this interface to its neighbors in their link-state PDUs.	**detail extensive**
Adjacency advertisement: Suppress	This neighbor has signaled not to advertise the interface in the routing device's outbound link-state PDUs.	**detail extensive**

Table 151: show isis adjacency Output Fields (continued)

Field Name	Field Description	Level of Output
IP addresses	IP address of this neighbor.	**detail extensive**
IPV6 Address	IPv6 address of the neighbor.	**detail extensive**
Level 1 IPv4 Adj-SID	Level 1 IPv4 node-SID of the adjacent neighbor.	**detail extensive**
Level 1 IPv6 Adj-SID	Level 1 IPv6 node-SID of the adjacent neighbor.	
Level 2 IPv4 Adj-SID	Level 2 IPv4 node-SID of the adjacent neighbor.	**detail extensive**
Level 2 IPv6 Adj-SID	Level 2 IPv6 node-SID of the adjacent neighbor.	**detail extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 151: show isis adjacency Output Fields (continued)

Field Name	Field Description	Level of Output

Transition log	List of recent transitions, including:
	• **When**—Time at which an IS-IS adjacency transition occurred.
	• **State**—Current state of the IS-IS adjacency (**up**, **down**, or **rejected**).
	  ◦ **Up**—Adjacency is up and operational.
	  ◦ **Down**—Adjacency is down and not available.
	  ◦ **Rejected**—Adjacency has been rejected.
	• **Event**—Type of transition that occurred.
	  ◦ **Seenself**—Possible routing loop has been detected.
	  ◦ **Interface down**—IS-IS interface has gone down and is no longer available.
	  ◦ **Error**—Adjacency error.
	• **Down reason**—Reason that an IS-IS adjacency is down:
	  ◦ **3-Way Handshake Failed**—Connection establishment failed.
	  ◦ **Address Mismatch**—Address mismatch caused link failure.
	  ◦ **Aged Out**—Link expired.
	  ◦ **ISO Area Mismatch**—IS-IS area mismatch caused link failure.
	  ◦ **Bad Hello**—Unacceptable hello message caused link failure.
	  ◦ **BFD Session Down**—Bidirectional failure detection caused link failure.
	  ◦ **Interface Disabled**—IS-IS interface is disabled.
	  ◦ **Interface Down**—IS-IS interface is unavailable.
	  ◦ **Interface Level Disabled**—IS-IS level is disabled.
	  ◦ **Level Changed**—IS-IS level has changed on the adjacency.
	  ◦ **Level Mismatch**—Levels on adjacency are not compatible.
	  ◦ **MPLS LSP Down**—Label-switched path (LSP) is unavailable.
	  ◦ **MT Topology Changed**—IS-IS topology has changed.
	  ◦ **MT Topology Mismatch**—IS-IS topology is mismatched.
	  ◦ **Remote System ID Changed**—Adjacency peer system ID changed.
	  ◦ **Protocol Shutdown**—IS-IS protocol is disabled.
	  ◦ **CLI Command**—Adjacency brought down by user.
	  ◦ **Unknown**—Unknown.
		extensive

Sample Output

show isis adjacency

```
user@host> show isis adjacency
Interface           System         L State        Hold (secs) SNPA
at-2/3/0.0          ranier         3  Up                   23
```

show isis adjacency brief

The output for the **show isis adjacency brief** command is identical to that for the **show isis adjacency** command. For sample output, see show isis adjacency on page 1983.

Copyright © 2018, Juniper Networks, Inc.

show isis adjacency detail

```
user@host> show isis adjacency detail
ranier
  Interface: at-2/3/0.0, Level: 3, State: Up, Expires in 21 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 00:01:09 ago
  Circuit type: 3, Speaks: IP, IPv6
  Topologies: Unicast, IPV6-Unicast  Restart capable: Yes, Adjacency advertisement:
 Advertise
  LAN id: pro-bng3-c-F.02, IP addresses: 11.1.1.2
  IPv6 addresses: fe80::2a0:a514:0:4745
  Level 1 IPv4 Adj-SID: 299808, IPv6 Adj-SID: 299824
```

show isis adjacency extensive

```
user@host> show isis adjacency extensive
ranier
  Interface: at-2/3/0.0, Level: 3, State: Up, Expires in 22 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 00:01:16 ago
  Circuit type: 3, Speaks: IP, IPv6

  Topologies: Unicast, IPV6-Unicast
  Restart capable: Yes, Adjacency advertisement: Advertise
  IP addresses: 11.1.1.2
  IPv6 addresses: fe80::2a0:a514:0:3e45
  Level 1 IPv4 Adj-SID: 300112, IPv6 Adj-SID: 300304
  Level 2 IPv4 Adj-SID: 300320, IPv6 Adj-SID: 300336
  Transition log:
  When                  State        Event           Down reason
  Thu Mar 26 06:13:18   Up           Seenself
```


Copyright © 2018, Juniper Networks, Inc.

show isis authentication

List of Syntax

Syntax on page 1985
Syntax (EX Series Switches and QFX Series) on page 1985

Syntax

show isis authentication
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches and QFX Series)

show isis authentication
<instance *instance-name*>

Release Information

Command introduced in Junos OS Release 7.5.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Support for hitless authentication key rollover introduced in Junos OS Release 11.2.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display information about Intermediate System-to-Intermediate System (IS-IS) authentication.

Options

none—Display information about IS-IS authentication.

instance *instance-name*—(Optional) Display IS-IS authentication for the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- *show security keychain*

List of Sample Output

show isis authentication on page 1986
show isis authentication (With Hitless Authentication Key Rollover Configured) on page 1986

Output Fields

Table 152 on page 1985 describes the output fields for the **show isis authentication** command. Output fields are listed in the approximate order in which they appear.

Table 152: show isis authentication Output Fields

Field Name	Field Description
Interface	Interface name.

Table 152: show isis authentication Output Fields (continued)

Field Name	Field Description
Level	IS-IS level.
IIH Auth	IS-IS Hello (IIH) packet authentication type. Displays the name of the active keychain if hitless authentication key rollover is configured.
CSN Auth	Complete sequence number authentication type.
PSN Auth	Partial sequence number authentication type.
L1 LSP Authentication	Layer 1 link-state PDU authentication type.
L2 LSP Authentication	Layer 2 link-state PDU authentication type.

Sample Output

show isis authentication

```
user@host> show isis authentication
Interface             Level IIH Auth   CSN Auth   PSN Auth
at-2/3/0.0            1     Simple     Simple     Simple
                      2     MD5        MD5        MD5

L1 LSP Authentication: Simple
L2 LSP Authentication: MD5
```

show isis authentication (With Hitless Authentication Key Rollover Configured)

```
user@host> show isis authentication
Interface             Level IIH Auth   CSN Auth   PSN Auth
so-0/1/3.0            2     hakrhello  MD5        MD5

L2 LSP Authentication: MD5
```

 Copyright © 2018, Juniper Networks, Inc.

show isis backup coverage

Syntax

show isis backup coverage
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches and QFX Series)

show isis backup coverage
<instance *instance-name*>

Release Information

Command introduced in Junos OS Release 9.5.
Command introduced in Junos OS Release 9.5 for EX Series switches.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display information about the level of backup coverage available.

Options

none—Display information about the level of backup coverage available for all the nodes and prefixes in the network.

instance *instance-name*—(Optional) Display information about the level of backup coverage for a specific IS-IS routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- *Understanding Loop-Free Alternate Routes for IS-IS*
- *Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN*
- show isis backup label-switched-path on page 1989

List of Sample Output

show isis backup coverage on page 1988

Output Fields

Table 153 on page 1987 lists the output fields for the **show isis backup coverage** command. Output fields are listed in the approximate order in which they appear.

Table 153: show isis backup coverage Output Fields

Field Name	Field Description
Topology	Type of topology or address family: **IPV4 Unicast** or **IPV6 Unicast**.

Copyright © 2018, Juniper Networks, Inc.

Table 153: show isis backup coverage Output Fields (continued)

Field Name	Field Description
Level	IS-IS level: • **1**—Level 1 • **2**—Level 2
Node	By topology, the percentage of all routes configured on the node that are protected through backup coverage.
IPv4	Percentage of IPv4 unicast routes that are protected through backup coverage.
IPv6	Percentage of IPv6 unicast routes that are protected through backup coverage.
CLNS	Percentage of Connectionless Network Service (CLNS) routes that are protected through backup coverage.

Sample Output

show isis backup coverage

```
user@host> show isis backup coverage
Backup Coverage:
    Topology        Level    Node    IPv4     IPv6     CLNS
    IPV4 Unicast        2  28.57%  22.22%    0.00%    0.00%
    IPV6 Unicast        2   0.00%   0.00%    0.00%    0.00%
```


Copyright © 2018, Juniper Networks, Inc.

show isis backup label-switched-path

Syntax

show isis backup label-switched-path
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches and QFX Series)

show isis backup label-switched-path

Release Information

Command introduced in Junos OS Release 9.5.
Command introduced in Junos OS Release 9.5 for EX Series switches.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display information about MPLS label-switched-paths (LSPs) designated as backup routes for IS-IS routes.

Options

none—Display information about MPLS LSPs designated as backup routes for IS-IS routes.

logical-system (all | *logical-system-name*—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- *Understanding Loop-Free Alternate Routes for IS-IS*
- *Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN*
- show isis backup coverage on page 1987

List of Sample Output

show isis backup label-switched-path on page 1990

Output Fields

Table 154 on page 1989 lists the output fields for the **show isis backup label-switched-path** command. Output fields are listed in the approximate order in which they appear.

Table 154: show isis backup label-switched-path Output Fields

Field Name	Field Description
Backup MPLS LSPs	List of MPLS LSPs designated as backup paths for IS-IS routes.
Egress	IP address of the egress routing device for the LSP.

Copyright © 2018, Juniper Networks, Inc.

Table 154: show isis backup label-switched-path Output Fields (continued)

Field Name	Field Description

Status	State of the LSP:
	• **Up**—The routing device can detect RSVP hello messages from the neighbor.
	• **Down**—The routing device has received one of the following indications:
	  ◦ Communication failure from the neighbor.
	  ◦ Communication from IGP that the neighbor is unavailable.
	  ◦ Change in the sequence numbers in the RSVP hello messages sent by the neighbor.
	• **Deleted**—LSP is no longer available as a backup path.

Last change	Time elapsed since the neighbor state changed either from up to down or from down to up. The format is *hh:mm:ss*.
TE-metric	Configured traffic engineering metric.
Metric	Configured metric.

Sample Output

show isis backup label-switched-path

```
user@host> show isis backup label-switched-path
 Backup MPLS LSPs:
 f-to-g, Egress: 192.168.1.4, Status: up, Last change: 06:12:03
    TE-metric: 9, Metric: 0
```

 Copyright © 2018, Juniper Networks, Inc.

show isis backup spf results

Syntax

show isis backup spf results
<instance *instance-name*>
<level (1 | 2)>
<logical-system (all | *logical-system-name*)>
<no-coverage>
<topology (ipv4-unicast | ipv6-multicast | ipv6-unicast | unicast)>

Syntax (EX Series Switches)

show isis backup spf results
<instance *instance-name*>
<level (1 | 2)>
<no-coverage>
<topology (ipv4-unicast | unicast)>

Release Information

Command introduced in Junos OS Release 9.5.

Description

Display information about IS-IS shortest-path-first (SPF) calculations for backup paths.

Options

none—Display information about IS-IS SPF calculations for all backup paths for all destination nodes.

instance *instance-name*—(Optional) Display SPF calculations for backup paths for the specified routing instance.

level (1 | 2)—(Optional) Display SPF calculations for the backup paths for the specified IS-IS level.

logical-system *logical-system-name*—(Optional) Display SPF calculations for the backup paths for all logical systems or on a particular logical system.

no-coverage—(Optional) Display SPF calculations only for destinations that do not have backup coverage.

topology (ipv4-multicast | ipv6-multicast | ipv6-unicast | unicast)—(Optional) Display SPF calculations for backup paths for the specified topology only.

Required Privilege Level

view

Related Documentation

- *Example: Configuring Link and Node Protection for IS-IS Routes*
- show isis backup coverage on page 1987
- *Understanding Loop-Free Alternate Routes for IS-IS*
- *Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN*

List of Sample Output

show isis backup spf results on page 1992 show isis backup spf results no-coverage on page 1993

Output Fields Table 155 on page 1992 lists the output fields for the **show isis backup spf results** command. Output fields are listed in the approximate order in which they appear.

Table 155: show isis backup spf results Output Fields

Field Name	Field Description
node-name	Name of the destination node.
Address	Address of the destination node.
Primary next-hop	Interface and name of the node of the primary next hop to reach the destination.
Root	Name of the next-hop neighbor.
Metric	Metric to the node.
Eligible	Indicates that the next-hop neighbor has been designated as a backup path to the destination node.
Backup next-hop	Name of the interface of the backup next hop.
SNPA	Subnetwork point of attachment (MAC address of the next hop).
LSP	Name of the MPLS label-switched path (LSP) designated as a backup path.
Not eligible	Indicates that the next-hop neighbor cannot function as a backup path to the destination.
Reason	Describes why the next-hop neighbor is designated as **Not eligible** as a backup path.

Sample Output

show isis backup spf results

```
user@host> show isis backup spf results

IS-IS level 1 SPF results:
  0 nodes

 IS-IS level 2 SPF results:
banff.00
  Primary next-hop: so-6/0/0.0, IPV4, olympic
  Primary next-hop: ae0.0, IPV4, camaro, SNPA:  0:90:69:f:67:f0
  Primary next-hop: so-6/0/0.0, IPV6, olympic
  Primary next-hop: ae0.0, IPV6, camaro, SNPA:  0:90:69:f:67:f0
    Root: camaro, Root Metric: 10, Metric: 10
      Not eligible, Reason: Primary next-hop multipath
    Root: olympic, Root Metric: 10, Metric: 10
```

 Copyright © 2018, Juniper Networks, Inc.

```
      Not eligible, Reason: Primary next-hop multipath
    Root: glacier, Root Metric: 10, Metric: 25
      Not eligible, Reason: Primary next-hop multipath
crater.00
  Primary next-hop: so-6/0/0.0, IPV4, olympic
  Primary next-hop: so-6/0/0.0, IPV6, olympic
    Root: olympic, Root Metric: 10, Metric: 10
      Not eligible, Reason: Primary next-hop link fate sharing
    Root: glacier, Root Metric: 10, Metric: 15
      Eligible, Backup next-hop: as0.0, IPV4, glacier
      Eligible, Backup next-hop: as0.0, IPV6, glacier
    Root: camaro, Root Metric: 10, Metric: 20
      Not eligible, Reason: Interface is already covered
olympic.00
  Primary next-hop: so-6/0/0.0, IPV4, olympic
  Primary next-hop: so-6/0/0.0, IPV6, olympic
    Root: olympic, Root Metric: 10, Metric: 0
      Not eligible, Reason: Primary next-hop link fate sharing
    Root: camaro, Root Metric: 10, Metric: 20
      track-item: olympic.00-00
      track-item: kobuk.00-00
      Not eligible, Reason: Path loops
    Root: glacier, Root Metric: 10, Metric: 20
      track-item: olympic.00-00
      track-item: kobuk.00-00
      Not eligible, Reason: Path loops
camaro.00
  Primary next-hop: ae0.0, IPV4, camaro, SNPA:  0:90:69:f:67:f0
  Primary next-hop: ae0.0, IPV6, camaro, SNPA:  0:90:69:f:67:f0
    Root: camaro, Root Metric: 10, Metric: 0
      Not eligible, Reason: Primary next-hop link fate sharing
    Root: glacier, Root Metric: 10, Metric: 20
      track-item: camaro.00-00
      track-item: kobuk.00-00
      Not eligible, Reason: Path loops
    Root: olympic, Root Metric: 10, Metric: 20
      track-item: camaro.00-00
      track-item: kobuk.00-00
      Not eligible, Reason: Path loops
glacier.00
  Primary next-hop: as0.0, IPV4, glacier
  Primary next-hop: as0.0, IPV6, glacier
    Root: glacier, Root Metric: 10, Metric: 0
      Not eligible, Reason: Primary next-hop link fate sharing
    Root: camaro, Root Metric: 10, Metric: 20
      track-item: glacier.00-00
      track-item: kobuk.00-00
      Not eligible, Reason: Path loops
    Root: olympic, Root Metric: 10, Metric: 20
      track-item: glacier.00-00
      track-item: kobuk.00-00
      Not eligible, Reason: Path loops
  5 nodes
```

show isis backup spf results no-coverage

```
user@host> show isis backup spf results no-coverage
IS-IS level 1 SPF results:
pro-bng3-k.00
  Primary next-hop: fe-1/3/3.0, IPV4, pro-bng3-k, SNPA:  b0:c6:9a:2c:f0:de
  Primary next-hop: fe-1/3/3.0, IPV6, pro-bng3-k, SNPA:  b0:c6:9a:2c:f0:de
```

```
    Root: pro-bng3-k, Root Metric: 10, Metric: 0, Root Preference: 0x0
    Root: pro-bng3-i, Root Metric: 10, Metric: 20, Root Preference: 0x0
      track-item: pro-bng3-k.00-00
      track-item: pro-bng3-j.00-00
pro-bng3-i.00
  Primary next-hop: fe-0/1/2.0, IPV4, pro-bng3-i, SNPA:  b0:c6:9a:2a:f4:21
  Primary next-hop: fe-0/1/2.0, IPV6, pro-bng3-i, SNPA:  b0:c6:9a:2a:f4:21
    Root: pro-bng3-i, Root Metric: 10, Metric: 0, Root Preference: 0x0
    Root: pro-bng3-k, Root Metric: 10, Metric: 20, Root Preference: 0x0
      track-item: pro-bng3-j.00-00
      track-item: pro-bng3-i.00-00
   2 nodes

 IS-IS level 2 SPF results:
olympic.00
  Primary next-hop: so-6/0/0.0, IPV4, olympic
  Primary next-hop: so-6/0/0.0, IPV6, olympic
    Root: olympic, Root Metric: 10, Metric: 0
      Not eligible, Reason: Primary next-hop link fate sharing
    Root: camaro, Root Metric: 10, Metric: 20
      track-item: olympic.00-00
      track-item: kobuk.00-00
      Not eligible, Reason: Path loops
    Root: glacier, Root Metric: 10, Metric: 20
      track-item: olympic.00-00
      track-item: kobuk.00-00
      Not eligible, Reason: Path loops
camaro.00
  Primary next-hop: ae0.0, IPV4, camaro, SNPA:  0:90:69:f:67:f0
  Primary next-hop: ae0.0, IPV6, camaro, SNPA:  0:90:69:f:67:f0
    Root: camaro, Root Metric: 10, Metric: 0
      Not eligible, Reason: Primary next-hop link fate sharing
    Root: glacier, Root Metric: 10, Metric: 20
      track-item: camaro.00-00
      track-item: kobuk.00-00
      Not eligible, Reason: Path loops
    Root: olympic, Root Metric: 10, Metric: 20
      track-item: camaro.00-00
      track-item: kobuk.00-00
      Not eligible, Reason: Path loops
glacier.00
  Primary next-hop: as0.0, IPV4, glacier
  Primary next-hop: as0.0, IPV6, glacier
    Root: glacier, Root Metric: 10, Metric: 0
      Not eligible, Reason: Primary next-hop link fate sharing
    Root: camaro, Root Metric: 10, Metric: 20
      track-item: glacier.00-00
      track-item: kobuk.00-00
      Not eligible, Reason: Path loops
    Root: olympic, Root Metric: 10, Metric: 20
      track-item: glacier.00-00
      track-item: kobuk.00-00
      Not eligible, Reason: Path loops
  3 nodes
```

 Copyright © 2018, Juniper Networks, Inc.

show isis context-identifier

Syntax

show isis context-identifier
<brief | detail | extensive>
<*identifier name*>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced in Junos OS Release 10.4.

Description

Display IS-IS context identifier information.

Options

brief | detail | extensive—(Optional) Display the specified level of output.

identifier name—(Optional) Display information about the specified context identifier.

instance *instance-name*—(Optional) Display entries for the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Display the context identifier information for all logical systems or for a particular logical system.

Required Privilege Level

View

Output Fields

Table 156 on page 1995 lists the output fields for the **show isis context-identifier** command. Output fields are listed in the approximate order in which they appear.

Table 156: show isis context-identifier Output Fields

Field Name	Field Description	Level of Output
Context	IPv4 address that defines a protection pair. The context is manually configured on both primary and protector PEs.	**detail**
Owner	Protocol that requires the context.	**detail**
Role	Role of the PE, which is either primary or protector.	**detail**
Primary	Name of the primary PE.	**detail**
Metric	Advertised interior gateway protocol (IGP) metric.	**detail**

Sample Output

```
user@host> show isis context-identifier detail

IS-IS context database:
Context                 Owner     Role       Primary      Metric
2.2.4.3                MPLS      Primary    pro3-e        1
```

```
Advertiser pro3-e, Router ID 10.255.245.198, Metric 1, Level 1
Advertiser pro3-e, Router ID 10.255.245.198, Metric 1, Level 2
Advertiser pro3-c, Router ID 10.255.245.196, Metric 11, Level 2
```

 Copyright © 2018, Juniper Networks, Inc.

show isis database

List of Syntax

Syntax on page 1997
Syntax (EX Series Switches and QFX Series) on page 1997

Syntax

show isis database
<system-id>
<brief | detail | extensive>
<instance *instance-name*>
<level (1 | 2)>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches and QFX Series)

show isis database
<system-id>
<brief | detail | extensive>
<level (1 | 2)>
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Command introduced in Junos OS Release 15.1X53-D30 for the QFX10002 switch.

Description

Display the entries in the Intermediate System-to-Intermediate System (IS-IS) link-state database, which contains data about PDU packets.

Options

none—Display standard information about IS-IS link-state database entries for all routing instances.

system id—(Optional) Display IS-IS link-state database entries for the specified intermediate system.

brief | detail | extensive—(Optional) Display the specified level of output.

instance *instance-name*—(Optional) Display IS-IS link-state database entries for the specified routing instance.

level (1 | 2)—(Optional) Display IS-IS link-state database entries for the specified IS-IS level.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- clear isis database on page 1974

List of Sample Output

show isis database on page 1999
show isis database brief on page 2000
show isis database detail on page 2000
show isis database extensive on page 2001
show isis database extensive (SPRING) on page 2010

Output Fields

Table 157 on page 1998 describes the output fields for the **show isis database** command. Output fields are listed in the approximate order in which they appear. Fields that contain internal IS-IS information useful only in troubleshooting obscure problems are not described in the table. For more details about these fields, contact your customer support representative.

Table 157: show isis database Output Fields

Field Name	Field Description	Level of Output
Interface name	Name of the interface on which the link-state PDU has been received; always **IS-IS** for this command.	All levels
level	Level of intermediate system: • **1**—Intermediate system routes within an area; when the destination is outside an area, it routes toward a Level 2 system. • **2**—Intermediate system routes between areas and toward other ASs.	All levels
LSP ID	Link-state PDU identifier.	All levels
Sequence	Sequence number of the link-state PDU.	All levels
Checksum	Checksum value of the link-state PDU.	All levels
Lifetime (secs)	Remaining lifetime of the link-state PDU, in seconds.	All levels
Attributes	Attributes of the specified database: **L1**, **L2**, **Overload**, or **Attached** (L1 only).	none **brief**
# LSPs	Total number of link-state PDUs in the specified link-state database.	none **brief**
IP prefix	Prefix advertised by this link-state PDU.	**detail extensive**
IS neighbor	IS-IS neighbor of the advertising system.	**detail extensive**
IP prefix	IPv4 prefix advertised by this link-state PDU.	**detail extensive**
V6 prefix	IPv6 prefix advertised by this link-state PDU.	**detail extensive**
Metric	Metric of the prefix or neighbor.	**detail extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 157: show isis database Output Fields (continued)

Field Name	Field Description	Level of Output
Header	• **LSP ID**—Link state PDU identifier of the header. • **Length**—Header length. • **Allocated Length**—Amount of length available for the header. • **Router ID**—Address of the local routing device. • **Remaining Lifetime**—Remaining lifetime of the link-state PDU, in seconds.	**extensive**
Packet	• **LSP ID**—The identifier for the link-state PDU. • **Length**—Packet length. • **Lifetime**—Remaining lifetime, in seconds. • **Checksum**—The checksum of the link-state PDU. • **Sequence**—The sequence number of the link-state PDU. Every time the link-state PDU is updated, this number increments. • **Attributes**—Packet attributes. • **NLPID**—Network layer protocol identifier. • **Fixed length**—Specifies the set length for the packet.	**extensive**
TLVs	• **Area Address**—Area addresses that the routing device can reach. • **Speaks**—Supported routing protocols. • **IP router id**—ID of the routing device (usually the IP address). • **IP address**—IPv4 address. • **Hostname**—Assigned name of the routing device. • **IP prefix**—IP prefix of the routing device. • **Metric**—IS-IS metric that measures the cost of the adjacency between the originating routing device and the advertised routing device. • **IP extended prefix**—Extended IP prefix of the routing device. • **IS neighbor**—Directly attached neighbor's name and metric. • **IS extended neighbor**—Directly attached neighbor's name, metric, IP address, local interface index, and remote interface index. The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205. • **Router Capability**—ID of the routing device and flag. NOTE: Router capability also specifies SPRING capability and SPRING algorithm when segment routing is enabled on the routing device.	**extensive**
Extended IS Reachability TLV	Type and length of the TLV is useful to identify the IS extended neighbors packed in this TLV.	**extensive**
SubTLV len	The length of the SubTLV of IS extended neighbor.	**extensive**

Sample Output

show isis database

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                      Sequence Checksum Lifetime Attributes
```

```
kobuk.00-00                      0x3   0x3167     1057 L1 L2
camaro.00-00                     0x5   0x770e     1091 L1 L2
ranier.00-00                     0x4   0xaa95     1091 L1 L2
glacier.00-00                    0x4   0x206f     1089 L1 L2
glacier.02-00                    0x1   0xd141     1089 L1 L2
badlands.00-00                   0x3   0x87a2     1093 L1 L2
  6 LSPs

IS-IS level 2 link-state database:
LSP ID                      Sequence Checksum Lifetime Attributes
kobuk.00-00                      0x6   0x8d6b     1096 L1 L2
camaro.00-00                     0x9   0x877b     1101 L1 L2
ranier.00-00                     0x8   0x855d     1103 L1 L2
glacier.00-00                    0x7   0xf892     1098 L1 L2
glacier.02-00                    0x1   0xd141     1089 L1 L2
badlands.00-00                   0x6    0x562     1105 L1 L2
  6 LSPs
```

show isis database brief

The output for the **show isis database brief** command is identical to that for the **show isis database** command. For sample output, see show isis database on page 1999.

show isis database detail

```
user@host> show isis database logical-system CE3 sisira.00-00 detail

IS-IS level 1 link-state database:

sisira.00-00 Sequence: 0x11, Checksum: 0x10fc, Lifetime: 975 secs
   IS neighbor: hemantha-CE3.02               Metric:       10
   ES neighbor: 0015.0015.0015                Metric:       10 Down
   ES neighbor: 0025.0025.0025                Metric:       10 Down
   ES neighbor: 0030.0030.0030                Metric:       10 Down
   ES neighbor: 0040.0040.0040                Metric:       10 Down
   ES neighbor: sisira                        Metric:        0
   IP prefix: 1.0.0.0/24                      Metric:       10 External Down
   IP prefix: 3.0.0.0/24                      Metric:       10 External Down
   IP prefix: 4.0.0.0/24                      Metric:       10 External Down
   IP prefix: 5.0.0.0/24                      Metric:       10 Internal Up
   IP prefix: 15.15.15.15/32                  Metric:       10 External Down
   IP prefix: 25.25.25.25/32                  Metric:       10 External Down
   IP prefix: 30.30.30.30/32                  Metric:       10 External Down
   IP prefix: 40.40.40.40/32                  Metric:       10 External Down
   IP prefix: 60.60.60.60/32                  Metric:        0 Internal Up

IS-IS level 2 link-state database:

sisira.00-00 Sequence: 0x13, Checksum: 0x69ac, Lifetime: 993 secs
   IS neighbor: hemantha-CE3.02               Metric:       10
   IP prefix: 1.0.0.0/24                      Metric:       10 External Down
   IP prefix: 3.0.0.0/24                      Metric:       10 External Down
   IP prefix: 4.0.0.0/24                      Metric:       10 External Down
   IP prefix: 5.0.0.0/24                      Metric:       10 Internal Up
   IP prefix: 15.15.15.15/32                  Metric:       10 External Down
   IP prefix: 25.25.25.25/32                  Metric:       10 External Down
   IP prefix: 30.30.30.30/32                  Metric:       10 External Down
   IP prefix: 40.40.40.40/32                  Metric:       10 External Down
   IP prefix: 50.50.50.50/32                  Metric:       10 Internal Up
   IP prefix: 60.60.60.60/32                  Metric:        0 Internal Up
```

 Copyright © 2018, Juniper Networks, Inc.

```
   ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0015.0015.0015/152
                                               Metric:       10 External Down
   ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0025.0025.0025/152
                                               Metric:       10 External Down
   ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0030.0030.0030/152
                                               Metric:       10 External Down
   ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0040.0040.0040/152
                                               Metric:       10 External Down
   ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0060.0060.0060/152
                                               Metric:        0 Internal Up
```

show isis database extensive

```
user@host> show isis database extensive

IS-IS level 1 link-state database:

sisira.00-00 Sequence: 0x11, Checksum: 0x10fc, Lifetime: 970 secs
   IS neighbor: hemantha-CE3.02                Metric:       10
     Two-way fragment: hemantha-CE3.02-00, Two-way first fragment:
hemantha-CE3.02-00
   ES neighbor: 0015.0015.0015                 Metric:       10 Down
   ES neighbor: 0025.0025.0025                 Metric:       10 Down
   ES neighbor: 0030.0030.0030                 Metric:       10 Down
   ES neighbor: 0040.0040.0040                 Metric:       10 Down
   ES neighbor: sisira                         Metric:        0
   IP prefix: 1.0.0.0/24                       Metric:       10 External Down
   IP prefix: 3.0.0.0/24                       Metric:       10 External Down
   IP prefix: 4.0.0.0/24                       Metric:       10 External Down
   IP prefix: 5.0.0.0/24                       Metric:       10 Internal Up
   IP prefix: 15.15.15.15/32                   Metric:       10 External Down
   IP prefix: 25.25.25.25/32                   Metric:       10 External Down
   IP prefix: 30.30.30.30/32                   Metric:       10 External Down
   IP prefix: 40.40.40.40/32                   Metric:       10 External Down
   IP prefix: 60.60.60.60/32                   Metric:        0 Internal Up

  Header: LSP ID: sisira.00-00, Length: 336 bytes
    Allocated length: 336 bytes, Router ID: 0.0.0.0
    Remaining lifetime: 970 secs, Level: 1, Interface: 333
    Estimated free bytes: 144, Actual free bytes: 0
    Aging timer expires in: 970 secs
    Protocols: IP, IPv6, CLNS

  Packet: LSP ID: sisira.00-00, Length: 336 bytes, Lifetime : 1198 secs
    Checksum: 0x10fc, Sequence: 0x11, Attributes: 0xb L1 L2 Attached
    NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
    Packet type: 18, Packet version: 1, Max area: 0

  TLVs:
    Area address: 60.0006.80ff.f800.0000.0108.0001 (13)
    Speaks: IP
    Speaks: IPV6
    Speaks: CLNP
    Hostname: sisira
    ES neighbor TLV: Internal, Metric: default 0, Up
      ES: sisira
    IS neighbor: hemantha-CE3.02, Internal, Metric: default 10
    IS extended neighbor: hemantha-CE3.02, Metric: default 10
    ES neighbor TLV: External, Metric: default 10, Down
      ES: 0040.0040.0040
    ES neighbor TLV: External, Metric: default 10, Down
```

```
      ES: 0025.0025.0025
    ES neighbor TLV: External, Metric: default 10, Down
      ES: 0015.0015.0015
    ES neighbor TLV: External, Metric: default 10, Down
      ES: 0030.0030.0030
    IP external prefix: 3.0.0.0/24, Internal, Metric: default 10, Down
    IP external prefix: 40.40.40.40/32, Internal, Metric: default 10, Down
    IP external prefix: 4.0.0.0/24, Internal, Metric: default 10, Down
    IP external prefix: 25.25.25.25/32, Internal, Metric: default 10, Down
    IP external prefix: 15.15.15.15/32, Internal, Metric: default 10, Down
    IP external prefix: 1.0.0.0/24, Internal, Metric: default 10, Down
    IP external prefix: 30.30.30.30/32, Internal, Metric: default 10, Down
    IP extended prefix: 3.0.0.0/24 metric 10 down
    IP extended prefix: 40.40.40.40/32 metric 10 down
    IP extended prefix: 4.0.0.0/24 metric 10 down
    IP extended prefix: 25.25.25.25/32 metric 10 down
    IP extended prefix: 15.15.15.15/32 metric 10 down
    IP extended prefix: 1.0.0.0/24 metric 10 down
    IP extended prefix: 30.30.30.30/32 metric 10 down
    IP prefix: 60.60.60.60/32, Internal, Metric: default 0, Up
    IP prefix: 5.0.0.0/24, Internal, Metric: default 10, Up
    IP extended prefix: 60.60.60.60/32 metric 0 up
    IP extended prefix: 5.0.0.0/24 metric 10 up
  No queued transmissions

IS-IS level 2 link-state database:

Router-A.00-00 Sequence: 0x5, Checksum: 0x3196, Lifetime: 1158 secs
   IS neighbor: Router-B.02                      Metric:       10
     Two-way fragment: Router-B.02-00, Two-way first fragment: Router-B.02-00
   IS neighbor: Router-E.02                      Metric:       10
     Two-way fragment: Router-E.02-00, Two-way first fragment: Router-E.02-00
   IP prefix: 10.0.0.0/30                     Metric:       10 Internal Up
   IP prefix: 10.0.0.4/30                     Metric:       10 Internal Up
   IP prefix: 192.168.0.1/32                  Metric:        0 Internal Up

  Header: LSP ID: Router-A.00-00, Length: 208 bytes
    Allocated length: 1492 bytes, Router ID: 192.168.0.1
    Remaining lifetime: 1158 secs, Level: 2, Interface: 0
    Estimated free bytes: 1233, Actual free bytes: 1284
    Aging timer expires in: 1158 secs
    Protocols: IP, IPv6

  Packet: LSP ID: Router-A.00-00, Length: 208 bytes, Lifetime : 1198 secs
    Checksum: 0x3196, Sequence: 0x5, Attributes: 0x3 <L1 L2>
    NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
    Packet type: 20, Packet version: 1, Max area: 0

  TLVs:
    Area address: 49.0002 (3)
    LSP Buffer Size: 1492
    Speaks: IP
    Speaks: IPV6
    IP router id: 192.168.0.1
    IP address: 192.168.0.1
    Hostname: Router-A
    IP prefix: 192.168.0.1/32, Internal, Metric: default 0, Up
    IP prefix: 10.0.0.4/30, Internal, Metric: default 10, Up
    IP prefix: 10.0.0.0/30, Internal, Metric: default 10, Up
    IP extended prefix: 192.168.0.1/32 metric 0 up
    IP extended prefix: 10.0.0.4/30 metric 10 up
```

 Copyright © 2018, Juniper Networks, Inc.

```
    IP extended prefix: 10.0.0.0/30 metric 10 up
    IS neighbor: Router-E.02, Internal, Metric: default 10
    IS neighbor: Router-B.02, Internal, Metric: default 10
    IS extended neighbor: Router-E.02, Metric: default 10
      IP address: 10.0.0.1
      Local interface index: 101, Remote interface index: 0
    IS extended neighbor: Router-B.02, Metric: default 10
      IP address: 10.0.0.5
      Local interface index: 102, Remote interface index: 0
  No queued transmissions

Router-B.00-00 Sequence: 0x5, Checksum: 0xf8f, Lifetime: 1183 secs
   IS neighbor: Router-B.02                      Metric:      10
     Two-way fragment: Router-B.02-00, Two-way first fragment: Router-B.02-00
   IS neighbor: Router-C.02                      Metric:      10
     Two-way fragment: Router-C.02-00, Two-way first fragment: Router-C.02-00
   IP prefix: 10.0.0.4/30                     Metric:      10 Internal Up
   IP prefix: 10.0.0.8/30                     Metric:      10 Internal Up
   IP prefix: 192.168.0.2/32                  Metric:       0 Internal Up

  Header: LSP ID: Router-B.00-00, Length: 208 bytes
    Allocated length: 284 bytes, Router ID: 192.168.0.2
    Remaining lifetime: 1183 secs, Level: 2, Interface: 102
    Estimated free bytes: 114, Actual free bytes: 76
    Aging timer expires in: 1183 secs
    Protocols: IP, IPv6

  Packet: LSP ID: Router-B.00-00, Length: 208 bytes, Lifetime : 1196 secs
    Checksum: 0xf8f, Sequence: 0x5, Attributes: 0x3 <L1 L2>
    NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
    Packet type: 20, Packet version: 1, Max area: 0

  TLVs:
    Area address: 49.0002 (3)
    LSP Buffer Size: 1492
    Speaks: IP
    Speaks: IPV6
    IP router id: 192.168.0.2
    IP address: 192.168.0.2
    Hostname: Router-B
    IP prefix: 192.168.0.2/32, Internal, Metric: default 0, Up
    IP prefix: 10.0.0.4/30, Internal, Metric: default 10, Up
    IP prefix: 10.0.0.8/30, Internal, Metric: default 10, Up
    IP extended prefix: 192.168.0.2/32 metric 0 up
    IP extended prefix: 10.0.0.4/30 metric 10 up
    IP extended prefix: 10.0.0.8/30 metric 10 up
    IS neighbor: Router-B.02, Internal, Metric: default 10
    IS neighbor: Router-C.02, Internal, Metric: default 10
    IS extended neighbor: Router-B.02, Metric: default 10
      IP address: 10.0.0.6
      Local interface index: 108, Remote interface index: 0
    IS extended neighbor: Router-C.02, Metric: default 10
      IP address: 10.0.0.9
      Local interface index: 109, Remote interface index: 0
  No queued transmissions

Router-B.02-00 Sequence: 0x1, Checksum: 0x3c7c, Lifetime: 1156 secs
   IS neighbor: Router-A.00                      Metric:       0
     Two-way fragment: Router-A.00-00, Two-way first fragment: Router-A.00-00
   IS neighbor: Router-B.00                      Metric:       0
     Two-way fragment: Router-B.00-00, Two-way first fragment: Router-B.00-00
```

Copyright © 2018, Juniper Networks, Inc.

```
  Header: LSP ID: Router-B.02-00, Length: 76 bytes
    Allocated length: 284 bytes, Router ID: 0.0.0.0
    Remaining lifetime: 1156 secs, Level: 2, Interface: 102
    Estimated free bytes: 208, Actual free bytes: 208
    Aging timer expires in: 1156 secs

  Packet: LSP ID: Router-B.02-00, Length: 76 bytes, Lifetime : 1196 secs
    Checksum: 0x3c7c, Sequence: 0x1, Attributes: 0x3 <L1 L2>
    NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
    Packet type: 20, Packet version: 1, Max area: 0

  TLVs:
    IS neighbor: Router-B.00, Internal, Metric: default 0
    IS neighbor: Router-A.00, Internal, Metric: default 0
    IS extended neighbor: Router-B.00, Metric: default 0
    IS extended neighbor: Router-A.00, Metric: default 0
  No queued transmissions

Router-C.00-00 Sequence: 0x5, Checksum: 0x255b, Lifetime: 1182 secs
   IS neighbor: Router-C.02                      Metric:      10
     Two-way fragment: Router-C.02-00, Two-way first fragment: Router-C.02-00
   IS neighbor: Router-D.03                      Metric:      10
     Two-way fragment: Router-D.03-00, Two-way first fragment: Router-D.03-00
   IP prefix: 10.0.0.8/30                     Metric:      10 Internal Up
   IP prefix: 10.0.0.12/30                    Metric:      10 Internal Up
   IP prefix: 192.168.0.3/32                  Metric:       0 Internal Up

  Header: LSP ID: Router-C.00-00, Length: 208 bytes
    Allocated length: 284 bytes, Router ID: 192.168.0.3
    Remaining lifetime: 1182 secs, Level: 2, Interface: 102
    Estimated free bytes: 114, Actual free bytes: 76
    Aging timer expires in: 1182 secs
    Protocols: IP, IPv6

  Packet: LSP ID: Router-C.00-00, Length: 208 bytes, Lifetime : 1196 secs
    Checksum: 0x255b, Sequence: 0x5, Attributes: 0x3 <L1 L2>
    NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
    Packet type: 20, Packet version: 1, Max area: 0

  TLVs:
    Area address: 49.0002 (3)
    LSP Buffer Size: 1492
    Speaks: IP
    Speaks: IPV6
    IP router id: 192.168.0.3
    IP address: 192.168.0.3
    Hostname: Router-C
    IP prefix: 192.168.0.3/32, Internal, Metric: default 0, Up
    IP prefix: 10.0.0.8/30, Internal, Metric: default 10, Up
    IP prefix: 10.0.0.12/30, Internal, Metric: default 10, Up
    IP extended prefix: 192.168.0.3/32 metric 0 up
    IP extended prefix: 10.0.0.8/30 metric 10 up
    IP extended prefix: 10.0.0.12/30 metric 10 up
    IS neighbor: Router-C.02, Internal, Metric: default 10
    IS neighbor: Router-D.03, Internal, Metric: default 10
    IS extended neighbor: Router-C.02, Metric: default 10
      IP address: 10.0.0.10
      Local interface index: 105, Remote interface index: 0
    IS extended neighbor: Router-D.03, Metric: default 10
      IP address: 10.0.0.13
```

 Copyright © 2018, Juniper Networks, Inc.

```
      Local interface index: 106, Remote interface index: 0
    No queued transmissions

  Router-C.02-00 Sequence: 0x1, Checksum: 0xaa09, Lifetime: 1181 secs
     IS neighbor: Router-B.00                      Metric:        0
       Two-way fragment: Router-B.00-00, Two-way first fragment: Router-B.00-00
     IS neighbor: Router-C.00                      Metric:        0
       Two-way fragment: Router-C.00-00, Two-way first fragment: Router-C.00-00

    Header: LSP ID: Router-C.02-00, Length: 76 bytes
      Allocated length: 284 bytes, Router ID: 0.0.0.0
      Remaining lifetime: 1181 secs, Level: 2, Interface: 102
      Estimated free bytes: 208, Actual free bytes: 208
      Aging timer expires in: 1181 secs

    Packet: LSP ID: Router-C.02-00, Length: 76 bytes, Lifetime : 1194 secs
      Checksum: 0xaa09, Sequence: 0x1, Attributes: 0x3 <L1 L2>
      NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
      Packet type: 20, Packet version: 1, Max area: 0

    TLVs:
      IS neighbor: Router-C.00, Internal, Metric: default 0
      IS neighbor: Router-B.00, Internal, Metric: default 0
      IS extended neighbor: Router-C.00, Metric: default 0
      IS extended neighbor: Router-B.00, Metric: default 0
    No queued transmissions

  Router-D.00-00 Sequence: 0x4, Checksum: 0x8ab7, Lifetime: 1180 secs
     IS neighbor: Router-D.02                      Metric:       10
       Two-way fragment: Router-D.02-00, Two-way first fragment: Router-D.02-00
     IS neighbor: Router-D.03                      Metric:       10
       Two-way fragment: Router-D.03-00, Two-way first fragment: Router-D.03-00
     IP prefix: 10.0.0.12/30                    Metric:       10 Internal Up
     IP prefix: 10.0.0.20/30                    Metric:       10 Internal Up
     IP prefix: 192.168.0.4/32                  Metric:        0 Internal Up

    Header: LSP ID: Router-D.00-00, Length: 208 bytes
      Allocated length: 284 bytes, Router ID: 192.168.0.4
      Remaining lifetime: 1180 secs, Level: 2, Interface: 102
      Estimated free bytes: 114, Actual free bytes: 76
      Aging timer expires in: 1180 secs
      Protocols: IP, IPv6

    Packet: LSP ID: Router-D.00-00, Length: 208 bytes, Lifetime : 1192 secs
      Checksum: 0x8ab7, Sequence: 0x4, Attributes: 0x3 <L1 L2>
      NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
      Packet type: 20, Packet version: 1, Max area: 0

    TLVs:
      Area address: 49.0002 (3)
      LSP Buffer Size: 1492
      Speaks: IP
      Speaks: IPV6
      IP router id: 192.168.0.4
      IP address: 192.168.0.4
      Hostname: Router-D
      IP prefix: 192.168.0.4/32, Internal, Metric: default 0, Up
      IP prefix: 10.0.0.12/30, Internal, Metric: default 10, Up
      IP prefix: 10.0.0.20/30, Internal, Metric: default 10, Up
      IP extended prefix: 192.168.0.4/32 metric 0 up
      IP extended prefix: 10.0.0.12/30 metric 10 up
```

```
    IP extended prefix: 10.0.0.20/30 metric 10 up
    IS neighbor: Router-D.02, Internal, Metric: default 10
    IS neighbor: Router-D.03, Internal, Metric: default 10
    IS extended neighbor: Router-D.02, Metric: default 10
      IP address: 10.0.0.22
      Local interface index: 115, Remote interface index: 0
    IS extended neighbor: Router-D.03, Metric: default 10
      IP address: 10.0.0.14
      Local interface index: 114, Remote interface index: 0
  No queued transmissions

Router-D.02-00 Sequence: 0x1, Checksum: 0xebbc, Lifetime: 1128 secs
   IS neighbor: Router-D.00                      Metric:       0
     Two-way fragment: Router-D.00-00, Two-way first fragment: Router-D.00-00
   IS neighbor: Router-F.00                      Metric:       0
     Two-way fragment: Router-F.00-00, Two-way first fragment: Router-F.00-00

  Header: LSP ID: Router-D.02-00, Length: 76 bytes
    Allocated length: 284 bytes, Router ID: 0.0.0.0
    Remaining lifetime: 1128 secs, Level: 2, Interface: 101
    Estimated free bytes: 208, Actual free bytes: 208
    Aging timer expires in: 1128 secs

  Packet: LSP ID: Router-D.02-00, Length: 76 bytes, Lifetime : 1160 secs
    Checksum: 0xebbc, Sequence: 0x1, Attributes: 0x3 <L1 L2>
    NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
    Packet type: 20, Packet version: 1, Max area: 0

  TLVs:
    IS neighbor: Router-D.00, Internal, Metric: default 0
    IS neighbor: Router-F.00, Internal, Metric: default 0
    IS extended neighbor: Router-D.00, Metric: default 0
    IS extended neighbor: Router-F.00, Metric: default 0
  No queued transmissions

Router-D.03-00 Sequence: 0x1, Checksum: 0x129b, Lifetime: 1180 secs
   IS neighbor: Router-C.00                      Metric:       0
     Two-way fragment: Router-C.00-00, Two-way first fragment: Router-C.00-00
   IS neighbor: Router-D.00                      Metric:       0
     Two-way fragment: Router-D.00-00, Two-way first fragment: Router-D.00-00

  Header: LSP ID: Router-D.03-00, Length: 76 bytes
    Allocated length: 284 bytes, Router ID: 0.0.0.0
    Remaining lifetime: 1180 secs, Level: 2, Interface: 101
    Estimated free bytes: 208, Actual free bytes: 208
    Aging timer expires in: 1180 secs

  Packet: LSP ID: Router-D.03-00, Length: 76 bytes, Lifetime : 1192 secs
    Checksum: 0x129b, Sequence: 0x1, Attributes: 0x3 <L1 L2>
    NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
    Packet type: 20, Packet version: 1, Max area: 0

  TLVs:
    IS neighbor: Router-D.00, Internal, Metric: default 0
    IS neighbor: Router-C.00, Internal, Metric: default 0
    IS extended neighbor: Router-D.00, Metric: default 0
    IS extended neighbor: Router-C.00, Metric: default 0
  No queued transmissions

Router-E.00-00 Sequence: 0x4, Checksum: 0x9da9, Lifetime: 1155 secs
   IS neighbor: Router-E.02                      Metric:      10
```

 Copyright © 2018, Juniper Networks, Inc.

```
    Two-way fragment: Router-E.02-00, Two-way first fragment: Router-E.02-00
  IS neighbor: Router-F.02                      Metric:       20
    Two-way fragment: Router-F.02-00, Two-way first fragment: Router-F.02-00
  IP prefix: 10.0.0.0/30                     Metric:       10 Internal Up
  IP prefix: 10.0.0.16/30                    Metric:       20 Internal Up
  IP prefix: 192.168.0.5/32                  Metric:        0 Internal Up

 Header: LSP ID: Router-E.00-00, Length: 208 bytes
   Allocated length: 284 bytes, Router ID: 192.168.0.5
   Remaining lifetime: 1155 secs, Level: 2, Interface: 101
   Estimated free bytes: 114, Actual free bytes: 76
   Aging timer expires in: 1155 secs
   Protocols: IP, IPv6

 Packet: LSP ID: Router-E.00-00, Length: 208 bytes, Lifetime : 1185 secs
   Checksum: 0x9da9, Sequence: 0x4, Attributes: 0x3 <L1 L2>
   NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
   Packet type: 20, Packet version: 1, Max area: 0

 TLVs:
   Area address: 49.0002 (3)
   LSP Buffer Size: 1492
   Speaks: IP
   Speaks: IPV6
   IP router id: 192.168.0.5
   IP address: 192.168.0.5
   Hostname: Router-E
   IP prefix: 192.168.0.5/32, Internal, Metric: default 0, Up
   IP prefix: 10.0.0.16/30, Internal, Metric: default 20, Up
   IP prefix: 10.0.0.0/30, Internal, Metric: default 10, Up
   IP extended prefix: 192.168.0.5/32 metric 0 up
   IP extended prefix: 10.0.0.16/30 metric 20 up
   IP extended prefix: 10.0.0.0/30 metric 10 up
   IS neighbor: Router-E.02, Internal, Metric: default 10
   IS neighbor: Router-F.02, Internal, Metric: default 20
   IS extended neighbor: Router-E.02, Metric: default 10
     IP address: 10.0.0.2
     Local interface index: 112, Remote interface index: 0
   IS extended neighbor: Router-F.02, Metric: default 20
     IP address: 10.0.0.17
     Local interface index: 111, Remote interface index: 0
 No queued transmissions

Router-E.02-00 Sequence: 0x1, Checksum: 0xb4fa, Lifetime: 1130 secs
  IS neighbor: Router-A.00                      Metric:        0
    Two-way fragment: Router-A.00-00, Two-way first fragment: Router-A.00-00
  IS neighbor: Router-E.00                      Metric:        0
    Two-way fragment: Router-E.00-00, Two-way first fragment: Router-E.00-00

 Header: LSP ID: Router-E.02-00, Length: 76 bytes
   Allocated length: 284 bytes, Router ID: 0.0.0.0
   Remaining lifetime: 1130 secs, Level: 2, Interface: 101
   Estimated free bytes: 208, Actual free bytes: 208
   Aging timer expires in: 1130 secs

 Packet: LSP ID: Router-E.02-00, Length: 76 bytes, Lifetime : 1161 secs
   Checksum: 0xb4fa, Sequence: 0x1, Attributes: 0x3 <L1 L2>
   NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
   Packet type: 20, Packet version: 1, Max area: 0

 TLVs:
```

Copyright © 2018, Juniper Networks, Inc.

```
    IS neighbor: Router-E.00, Internal, Metric: default 0
    IS neighbor: Router-A.00, Internal, Metric: default 0
    IS extended neighbor: Router-E.00, Metric: default 0
    IS extended neighbor: Router-A.00, Metric: default 0
  No queued transmissions

Router-F.00-00 Sequence: 0x5, Checksum: 0x94bd, Lifetime: 1153 secs
   IS neighbor: Router-D.02                      Metric:      10
     Two-way fragment: Router-D.02-00, Two-way first fragment: Router-D.02-00
   IS neighbor: Router-F.02                      Metric:      10
     Two-way fragment: Router-F.02-00, Two-way first fragment: Router-F.02-00
   IP prefix: 10.0.0.16/30                    Metric:      10 Internal Up
   IP prefix: 10.0.0.20/30                    Metric:      10 Internal Up
   IP prefix: 192.168.0.6/32                  Metric:       0 Internal Up

  Header: LSP ID: Router-F.00-00, Length: 208 bytes
    Allocated length: 284 bytes, Router ID: 192.168.0.6
    Remaining lifetime: 1153 secs, Level: 2, Interface: 101
    Estimated free bytes: 76, Actual free bytes: 76
    Aging timer expires in: 1153 secs
    Protocols: IP, IPv6

  Packet: LSP ID: Router-F.00-00, Length: 208 bytes, Lifetime : 1183 secs
    Checksum: 0x94bd, Sequence: 0x5, Attributes: 0x3 <L1 L2>
    NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
    Packet type: 20, Packet version: 1, Max area: 0

  TLVs:
    Area address: 49.0002 (3)
    LSP Buffer Size: 1492
    Speaks: IP
    Speaks: IPV6
    IP router id: 192.168.0.6
    IP address: 192.168.0.6
    Hostname: Router-F
    IP prefix: 192.168.0.6/32, Internal, Metric: default 0, Up
    IP prefix: 10.0.0.16/30, Internal, Metric: default 10, Up
    IP prefix: 10.0.0.20/30, Internal, Metric: default 10, Up
    IP extended prefix: 192.168.0.6/32 metric 0 up
    IP extended prefix: 10.0.0.16/30 metric 10 up
    IP extended prefix: 10.0.0.20/30 metric 10 up
    IS neighbor: Router-D.02, Internal, Metric: default 10
    IS neighbor: Router-F.02, Internal, Metric: default 10
    IS extended neighbor: Router-D.02, Metric: default 10
      IP address: 10.0.0.21
      Local interface index: 94, Remote interface index: 0
    IS extended neighbor: Router-F.02, Metric: default 10
      IP address: 10.0.0.18
      Local interface index: 93, Remote interface index: 0
  No queued transmissions

Router-E.02-00 Sequence: 0x1, Checksum: 0xb4fa, Lifetime: 1130 secs
   IS neighbor: Router-A.00                      Metric:       0
     Two-way fragment: Router-A.00-00, Two-way first fragment: Router-A.00-00
   IS neighbor: Router-E.00                      Metric:       0
     Two-way fragment: Router-E.00-00, Two-way first fragment: Router-E.00-00

  Header: LSP ID: Router-E.02-00, Length: 76 bytes
    Allocated length: 284 bytes, Router ID: 0.0.0.0
    Remaining lifetime: 1130 secs, Level: 2, Interface: 101
    Estimated free bytes: 208, Actual free bytes: 208
```

 Copyright © 2018, Juniper Networks, Inc.

```
    Aging timer expires in: 1130 secs

  Packet: LSP ID: Router-E.02-00, Length: 76 bytes, Lifetime : 1161 secs
    Checksum: 0xb4fa, Sequence: 0x1, Attributes: 0x3 <L1 L2>
    NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
    Packet type: 20, Packet version: 1, Max area: 0

  TLVs:
    IS neighbor: Router-E.00, Internal, Metric: default 0
    IS neighbor: Router-A.00, Internal, Metric: default 0
    IS extended neighbor: Router-E.00, Metric: default 0
    IS extended neighbor: Router-A.00, Metric: default 0
  No queued transmissions

Router-F.00-00 Sequence: 0x5, Checksum: 0x94bd, Lifetime: 1153 secs
   IS neighbor: Router-D.02                     Metric:      10
     Two-way fragment: Router-D.02-00, Two-way first fragment: Router-D.02-00
   IS neighbor: Router-F.02                     Metric:      10
     Two-way fragment: Router-F.02-00, Two-way first fragment: Router-F.02-00
   IP prefix: 10.0.0.16/30                   Metric:      10 Internal Up
   IP prefix: 10.0.0.20/30                   Metric:      10 Internal Up
   IP prefix: 192.168.0.6/32                 Metric:       0 Internal Up

  Header: LSP ID: Router-F.00-00, Length: 208 bytes
    Allocated length: 284 bytes, Router ID: 192.168.0.6
    Remaining lifetime: 1153 secs, Level: 2, Interface: 101
    Estimated free bytes: 76, Actual free bytes: 76
    Aging timer expires in: 1153 secs
    Protocols: IP, IPv6

  Packet: LSP ID: Router-F.00-00, Length: 208 bytes, Lifetime : 1183 secs
    Checksum: 0x94bd, Sequence: 0x5, Attributes: 0x3 <L1 L2>
    NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
    Packet type: 20, Packet version: 1, Max area: 0

  TLVs:
    Area address: 49.0002 (3)
    LSP Buffer Size: 1492
    Speaks: IP
    Speaks: IPV6
    IP router id: 192.168.0.6
    IP address: 192.168.0.6
    Hostname: Router-F
    IP prefix: 192.168.0.6/32, Internal, Metric: default 0, Up
    IP prefix: 10.0.0.16/30, Internal, Metric: default 10, Up
    IP prefix: 10.0.0.20/30, Internal, Metric: default 10, Up
    IP extended prefix: 192.168.0.6/32 metric 0 up
    IP extended prefix: 10.0.0.16/30 metric 10 up
    IP extended prefix: 10.0.0.20/30 metric 10 up
    IS neighbor: Router-D.02, Internal, Metric: default 10
    IS neighbor: Router-F.02, Internal, Metric: default 10
    IS extended neighbor: Router-D.02, Metric: default 10
      IP address: 10.0.0.21
      Local interface index: 94, Remote interface index: 0
    IS extended neighbor: Router-F.02, Metric: default 10
      IP address: 10.0.0.18
      Local interface index: 93, Remote interface index: 0
  No queued transmissions

Router-F.02-00 Sequence: 0x1, Checksum: 0xf5ae, Lifetime: 1153 secs
   IS neighbor: Router-E.00                     Metric:       0
```

```
    Two-way fragment: Router-E.00-00, Two-way first fragment: Router-E.00-00
  IS neighbor: Router-F.00                      Metric:        0
    Two-way fragment: Router-F.00-00, Two-way first fragment: Router-F.00-00

 Header: LSP ID: Router-F.02-00, Length: 76 bytes
   Allocated length: 284 bytes, Router ID: 0.0.0.0
   Remaining lifetime: 1153 secs, Level: 2, Interface: 101
   Estimated free bytes: 208, Actual free bytes: 208
   Aging timer expires in: 1153 secs

 Packet: LSP ID: Router-F.02-00, Length: 76 bytes, Lifetime : 1183 secs
   Checksum: 0xf5ae, Sequence: 0x1, Attributes: 0x3 <L1 L2>
   NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
   Packet type: 20, Packet version: 1, Max area: 0

 TLVs:
   IS neighbor: Router-F.00, Internal, Metric: default 0
   IS neighbor: Router-E.00, Internal, Metric: default 0
   IS extended neighbor: Router-F.00, Metric: default 0
   IS extended neighbor: Router-E.00, Metric: default 0
 No queued transmissions
```

show isis database extensive (SPRING)

```
user@host> show isis database extensive
Extended IS Reachability TLV, Type: 22, Length: 99
    IS extended neighbor: r0.00, Metric: default 10 SubTLV len: 88
      IP address: 11.1.1.2
      Neighbor's IP address: 11.1.1.1
      Local interface index: 333, Remote interface index: 342
      Current reservable bandwidth:
        Priority 0 : 1000Mbps
        Priority 1 : 1000Mbps
        Priority 2 : 1000Mbps
        Priority 3 : 1000Mbps
        Priority 4 : 1000Mbps
        Priority 5 : 1000Mbps
        Priority 6 : 1000Mbps
        Priority 7 : 1000Mbps
Maximum reservable bandwidth: 1000Mbps
      Maximum bandwidth: 1000Mbps
      Administrative groups:  0 <none>
P2P IPV6 Adj-SID - Flags:0xf0(F:1,B:1,V:1,L:1,S:0,P:0), Weight:0, Label: 401
      P2P IPv6 Adj-SID:     401, Weight:   0, Flags: FBVL--
      P2P IPV4 Adj-SID - Flags:0x70(F:0,B:1,V:1,L:1,S:0,P:0), Weight:0, Label:
400
      P2P IPv4 Adj-SID:     400, Weight:   0, Flags: -BVL--
    Extended IS Reachability TLV, Type: 22, Length: 144
    IS extended neighbor: r1.06, Metric: default 10 SubTLV len: 133
      IP address: 12.1.1.1
      Local interface index: 336, Remote interface index: 0
      Current reservable bandwidth:
        Priority 0 : 1000Mbps
        Priority 1 : 1000Mbps
        Priority 2 : 1000Mbps
        Priority 3 : 1000Mbps
        Priority 4 : 1000Mbps
        Priority 5 : 1000Mbps
        Priority 6 : 1000Mbps
        Priority 7 : 1000Mbps
      Maximum reservable bandwidth: 1000Mbps
```

 Copyright © 2018, Juniper Networks, Inc.

```
      Maximum bandwidth: 1000Mbps
    Administrative groups:  0 <none>
LAN IPV6 Adj-SID -, Flags:0xfc(F:1,B:1,V:1,L:1,S:1,P:1), Weight:1
         Neighbor:r2, Label:803125
      LAN IPv6 Adj-SID:  803125, Weight:   1, Neighbor: r2, Flags: FBVLSP
      LAN IPV4 Adj-SID -, Flags:0x3c(F:0,B:0,V:1,L:1,S:1,P:1), Weight:1
         Neighbor:r2, Label:803122
      LAN IPv4 Adj-SID:  803122, Weight:   1, Neighbor: r2, Flags: --VLSP
      LAN IPV4 Adj-SID -, Flags:0x7c(F:0,B:1,V:1,L:1,S:1,P:1), Weight:1
         Neighbor:r2, Label:803123
      LAN IPv4 Adj-SID:  803123, Weight:   1, Neighbor: r2, Flags: -BVLSP
      LAN IPV6 Adj-SID -, Flags:0xf0(F:1,B:1,V:1,L:1,S:0,P:0), Weight:0
         Neighbor:r2, Label:37
      LAN IPv6 Adj-SID:      37, Weight:   0, Neighbor: r2, Flags: FBVL--
      LAN IPV4 Adj-SID -, Flags:0x70(F:0,B:1,V:1,L:1,S:0,P:0), Weight:0
         Neighbor:r2, Label:36
      LAN IPv4 Adj-SID:      36, Weight:   0, Neighbor: r2, Flags: -BVL--
```

show isis hostname

List of Syntax

Syntax on page 2012
Syntax (EX Series Switches and QFX Series) on page 2012

Syntax

show isis hostname
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches and QFX Series)

show isis hostname

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display IS-IS hostname database information.

This command displays the system ID-to-name cache. The output shows if the mapping has been learned by receipt of a Hostname TLV #137 (type dynamic) configured in Junos OS with the **set system host-name** command, or a static mapping defined in Junos OS with the **set system static-host-mapping *hostname* sysid** command (type static). The local router always has its type set to static even if **static-host-mapping** is not configured.

Options

none—Display IS-IS hostname database information.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show isis hostname on page 2013

Output Fields

Table 158 on page 2012 describes the output fields for the **show isis hostname** command. Output fields are listed in the approximate order in which they appear.

Table 158: show isis hostname Output Fields

Field Name	Field Description
System Id	System identifier mapped to the hostname.
Hostname	Hostname mapped to the system identifier.

 Copyright © 2018, Juniper Networks, Inc.

Table 158: show isis hostname Output Fields (continued)

Field Name	Field Description
Type	Type of mapping between system identifier and hostname. • **Dynamic**—Hostname mapping determined as described in RFC 2763, *Dynamic Hostname Exchange Mechanism for IS-IS*. • **Static**—Hostname mapping configured by user.

Sample Output

show isis hostname

```
user@host> show isis hostname
   IS-IS hostname database:
   System Id      Hostname                                       Type
   1921.6800.4201 isis1                                          Dynamic
   1921.6800.4202 isis2                                          Static
   1921.6800.4203 isis3                                          Dynamic
```

Copyright © 2018, Juniper Networks, Inc.

show isis interface

List of Syntax

Syntax on page 2014
Syntax (EX Series Switches and QFX Series) on page 2014

Syntax

```
show isis interface
<brief | detail | extensive>
<interface-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches and QFX Series)

```
show isis interface
<brief | detail | extensive>
<interface-name>
```

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 12.1 for the QFX Series.

Description

Display status information about Intermediate System-to-Intermediate System (IS-IS)-enabled interfaces.

NOTE: If the configured metric for an IS-IS level is above 63, and the **wide-metrics-only** statement is not configured, the **show isis interface detail** command and the **show isis interface extensive** command display 63 as the metric value for that level. Configure the **wide-metrics-only** statement to generate metric values greater than 63 on a per IS-IS level basis.

The **show isis interface** command displays the configured metric value for an IS-IS level irrespective of whether is configured or not.

Options

none—Display standard information about all IS-IS-enabled interfaces.

brief | detail | extensive—(Optional) Display the specified level of output.

interface-name—(Optional) Display information about the specified interface only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- *Understanding Wide IS-IS Metrics for Traffic Engineering*
- *Example: Enabling Wide IS-IS Metrics for Traffic Engineering*

 Copyright © 2018, Juniper Networks, Inc.

List of Sample Output

show isis interface on page 2017
show isis interface brief on page 2017
show isis interface detail on page 2017
show isis interface extensive on page 2017

Output Fields

Table 159 on page 2015 describes the output fields for the **show isis interface** command. Output fields are listed in the approximate order in which they appear.

Table 159: show isis interface Output Fields

Field Name	Field Description	Level of Output
interface-name	Name of the interface.	**detail**
Designated router	Routing device selected by other routers that is responsible for sending link-state advertisements that describe the network. Used only on broadcast networks.	**detail**
Index	Interface index assigned by the Junos OS kernel.	**detail**
State	Internal implementation information.	**detail**
Circuit id	Circuit identifier. NOTE: Each IS-IS interface is assigned a circuit ID value to identify the interface within the linkstate database. All interfaces (loopback, broadcast, and so on) and all point-to-point links share the locally significant value of 0x01, and this value is not incremented.	**detail**
Circuit type	Circuit type: • **1**—Level 1 only • **2**—Level 2 only • **3**—Level 1 and Level 2	**detail**
LSP interval	Interval between link-state PDUs sent from the interface.	**detail**
CSNP interval	Interval between complete sequence number PDUs sent from the interface.	**detail extensive**
Sysid	System identifier.	**detail**
Interface	Interface through which the adjacency is made.	none **brief**
L or Level	Level: • **1**—Level 1 only • **2**—Level 2 only • **3**—Level 1 and Level 2 NOTE: The default IS-IS level on loopback interfaces are always same as the IS-IS level configured on other IS-IS interfaces in a router. You can also configure IS-IS level on loopback interfaces per your requirement.	All levels
CirID	Circuit identifier.	none **brief**

Copyright © 2018, Juniper Networks, Inc.

Table 159: show isis interface Output Fields (continued)

Field Name	Field Description	Level of Output
Level 1 DR	Level 1 designated intermediate system.	none **brief**
Level 2 DR	Level 2 designated intermediate system.	none **brief**
L1/L2 Metric	Interface's metric for Level 1 and Level 2. If there is no information, the metric is **0**.	none **brief**
Flood-group Area-ID	Flood-group is configured on a specific IS-IS interface NOTE: Seen only when Flood-group is configured.	**detail extensive**
Adjacency advertisement: Advertise	This routing device has signaled to advertise this interface to its neighbors in their label-switched paths (LSPs).	**detail extensive**
Adjacency advertisement: Suppress	This neighbor has signaled not to advertise this interface in the routing device's outbound LSPs.	**detail extensive**
Adjacencies	Number of adjacencies established on this interface.	**detail**
Priority	Priority value for this interface.	**detail**
Metric	Metric value for this interface.	**detail**
Hello(s) / Hello Interval	Interface's hello interval.	**detail extensive**
Hold(s) / Hold Time	Interface's hold time.	**detail extensive**
Designated Router	Router responsible for sending network link-state advertisements, which describe all the routing devices attached to the network.	**detail**
Hello padding	Type of hello padding: • **Adaptive**—On point-to-point connections, the hello packets are padded from the initial detection of a new neighbor until the neighbor verifies the adjacency as Up in the adjacency state TLV. If the neighbor does not support the adjacency state TLV, then padding continues. On LAN connections, padding starts from the initial detection of a new neighbor until there is at least one active adjacency on the interface. • **Loose**—(Default) The hello packet is padded from the initial detection of a new neighbor until the adjacency transitions to the Up state. • **Strict**—Padding is performed on all interface types and for all adjacency states, and is continuous.	**extensive**
LDP sync state	Current LDP synchronization state: **in sync**, **in holddown**, or **not supported**.	**extensive**
reason	Reason for being in the LDP sync state.	**extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 159: show isis interface Output Fields (continued)

Field Name	Field Description	Level of Output
config holdtime	Configured value of the hold timer.	**extensive**
remaining	If the state is not in sync and the hold time is not infinity, then this field displays the remaining hold time in seconds.	**extensive**
IIH max size	Configured value of IS-IS hello packets	extensive

Sample Output

show isis interface

```
user@host> show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR        Level 2 DR        L1/L2 Metric
at-2/3/0.0            3   0x1 Point to Point    Point to Point         10/10
lo0.0                 3   0x1 Passive           Passive                 0/0
```

show isis interface brief

The output for the **show isis interface brief** command is identical to that for the **show isis interface** command. For sample output, see show isis interface on page 2017.

show isis interface detail

```
user@host> show isis interface detail
IS-IS interface database:
at-2/3/0.0
  Index: 66, State: 0x6, Circuit id: 0x1, Circuit type: 3
  LSP interval: 100 ms, CSNP interval: 5 s
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    1             1       64     10     9.000       27
    2             1       64     10     9.000       27
lo0.0
  Index: 64, State: 0x6, Circuit id: 0x1, Circuit type: 0
  LSP interval: 100 ms, CSNP interval: disabled
  Adjacency advertisement: Advertise
  Protection Type: Node Link, No eligible Backup
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    1             0       64      0 Passive
    2             0       64      0 Passive
```

show isis interface extensive

```
user@host> show isis interface extensive
IS-IS interface database:
xe-6/1/0.0
  Index: 75, State: 0x6, Circuit id: 0x1, Circuit type: 2
  LSP interval: 100 ms, CSNP interval: 10 s, Loose Hello padding, IIH max size:
1505
  Adjacency advertisement: Advertise

  Flood-group Area-ID: 49.0001
```

Copyright © 2018, Juniper Networks, Inc.

```
Level 1
  Adjacencies: 0, Priority: 64, Metric: 10
  Disabled
Level 2
  Adjacencies: 1, Priority: 64, Metric: 10
  Hello Interval: 20.000 s, Hold Time: 60 s
  Designated Router: nemean.03
```

 Copyright © 2018, Juniper Networks, Inc.

show isis overview

Syntax

show isis overview
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches and QFX Series)

show isis overview
<instance *instance-name*>

Release Information

Command introduced in Junos OS Release 8.5.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display IS-IS overview information.

Options

none—Display standard overview information about IS-IS for all routing instances.

instance *instance-name*—(Optional) Display overview information for the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show isis overview on page 2021

Output Fields

Table 160 on page 2019 lists the output fields for the **show isis overview** command. Output fields are listed in the approximate order in which they appear.

Table 160: show isis overview Output Fields

Field Name	Field Description
Hostname	Name of the router.
Sysid	Part of the ISO address of the routing device.
Areaid	The area number of the routing device.
Instance	IS-IS routing instance.
Router ID	Router ID of the routing device.
Adjacency holddown	Adjacency holddown capability: **enabled** or **disabled**.

Copyright © 2018, Juniper Networks, Inc.

Table 160: show isis overview Output Fields (continued)

Field Name	Field Description
Maximum Areas	Maximum number of IS-IS areas advertised by the routing device.
LSP life time	Lifetime of the link-state PDU, in seconds.
Filter low life time LSPs up to	LSPs with a lifetime lower than this value are filtered out.
Attached bit evaluation	Attached bit capability: **enabled** or **disabled**.
SPF delay	Delay before performing consecutive shortest-path-first (SPF) calculations.
SPF holddown	Delay before performing additional SPF calculations after the maximum number of consecutive SPF calculations is reached.
SPF rapid runs	Maximum number of SPF calculations that can be performed in succession before the holddown timer begins.
Overload bit at startup is set	Overload bit capability is enabled.
Overload high metrics	Overload high metrics capability: **enabled** or **disabled**.
Overload timeout	Time period after which overload is reset and the time that remains before the timer is set to expire.
Traffic engineering	Traffic engineering capability: **enabled** or **disabled**.
Restart	Graceful restart capability: **enabled** or **disabled**.
Restart duration	Time period for complete reacquisition of IS-IS neighbors.
Helper mode	Graceful restart helper capability: **enabled** or **disabled**.
Level	IS-IS level: • **1**—Level 1 information • **2**—Level 2 information
IPv4 is enabled	IP Protocol version 4 capability is enabled.
IPv6 is enabled	IP Protocol version 6 capability is enabled.
Micro-loop avoidance	Micro-loop avoidance is enabled. Generally adjacent nodes converge faster than neighboring nodes causing traffic to loop. A route convergence delay is configured to avoid such micro loops.
Internal route preference	Preference value of internal routes.

 Copyright © 2018, Juniper Networks, Inc.

Table 160: show isis overview Output Fields (continued)

Field Name	Field Description
External route preference	Preference value of external routes.
Prefix export limit	Number of prefixes allowed to be exported, as configured by the **prefix-export-limit** statement.
Prefix export count	Number of prefixes exported.
Wide area metrics are enabled	Wide area metrics capability is enabled.
Narrow metrics are enabled	Narrow metrics capability is enabled.
Adjacency holddown is active	IS-IS adjacencies come up one after another when adjacency holddown is enabled.

Sample Output

show isis overview

```
user@host> show isis overview
Instance: master
  Router ID: 10.255.107.183
  Hostname: pro-bng3-a
  Sysid: 0192.0168.0001
  Areaid: 49.0002
  Adjacency holddown: enabled
  Maximum Areas: 3
  LSP life time: 1200
  Filter low life time LSPs up to: 300
  Attached bit evaluation: enabled
  SPF delay: 200 msec, SPF holddown: 5000 msec, SPF rapid runs: 3
  IPv4 is enabled, IPv6 is enabled
Micro-loop avoidance: Enabled
    Method: Route Convergence Delay, Route convergence delay: 5000 msec
  Traffic engineering: enabled
  Restart: Disabled
    Helper mode: Enabled
  Level 1
    Internal route preference: 15
    External route preference: 160
    Wide metrics are enabled, Narrow metrics are enabled
    Adjacency holddown is active
  Level 2
    Internal route preference: 18
    External route preference: 165
    Prefix export limit: 5, Prefix export count: 5
    Wide metrics are enabled
    Adjacency holddown is active


user@host> show isis overview logical-system R2
```

```
Instance: master
  Router ID: 192.168.0.2
  Hostname: pro-bng3-a-R2
  Sysid: 0192.0168.0002
  Areaid: 49.0002
  Adjacency holddown: enabled
  Maximum Areas: 3
  LSP life time: 1200
  Attached bit evaluation: enabled
  SPF delay: 200 msec, SPF holddown: 5000 msec, SPF rapid runs: 3
  IPv4 is enabled, IPv6 is enabled
  Traffic engineering: enabled
  Restart: Disabled
    Helper mode: Enabled
  Level 1
    Internal route preference: 15
    External route preference: 160
    Prefix export count: 0
    Wide metrics are enabled, Narrow metrics are enabled
  Level 2
    Internal route preference: 18
    External route preference: 165
    Prefix export count: 0
    Wide metrics are enabled, Narrow metrics are enabled
```

```
user@host> show isis overview logical-system R3
Instance: master
  Router ID: 192.168.0.3
  Hostname: pro-bng3-a-R3
  Sysid: 0192.0168.0003
  Areaid: 49.0002
  Adjacency holddown: enabled
  Maximum Areas: 3
  LSP life time: 1200
  Attached bit evaluation: enabled
  SPF delay: 200 msec, SPF holddown: 5000 msec, SPF rapid runs: 3
  IPv4 is enabled, IPv6 is enabled
  Traffic engineering: enabled
  Restart: Disabled
    Helper mode: Enabled
  Level 1
    Internal route preference: 15
    External route preference: 160
    Prefix export count: 0
    Wide metrics are enabled, Narrow metrics are enabled
  Level 2
    Internal route preference: 18
    External route preference: 165
    Prefix export count: 0
    Wide metrics are enabled, Narrow metrics are enabled
```

 Copyright © 2018, Juniper Networks, Inc.

show isis route

List of Syntax

Syntax on page 2023
Syntax (EX Series Switches and QFX Series) on page 2023

Syntax

show isis route
<*destination*>
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<topology (ipv4-multicast | ipv6-multicast | ipv6-unicast | unicast)>

Syntax (EX Series Switches and QFX Series)

show isis route
<*destination*>
<inet | inet6>
<instance *instance-name*>
<topology (ipv4-multicast | ipv6-multicast | ipv6-unicast | unicast)>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display the routes in the IS-IS routing table.

Options

none—Display all routes in the IS-IS routing table for all supported address families for all routing instances.

destination—(Optional) Destination address for the route.

inet | inet6—(Optional) Display inet (IPv4) or inet6 (IPv6) routes, respectively.

instance *instance-name*—(Optional) Display routes for the specified routing instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

topology (ipv4-multicast | ipv6-multicast | ipv6-unicast | unicast)—(Optional) Display routes for the specified topology only, or use unicast to display information, if available, for both IPv4 and IPv6 unicast topologies.

Required Privilege Level

view

List of Sample Output

show isis route logical-system on page 2024
show isis route floodgroup on page 2024
show isis route (CLNS) on page 2025
show isis route on page 2025

Output Fields Table 161 on page 2024 describes the output fields for the **show isis route** command. Output fields are listed in the approximate order in which they appear.

Table 161: show isis route Output Fields

Field Name	Field Description
Current version	Number of the current version of the IS-IS routing table.
L1	Version of Level 1 SPF that was run.
L2	Version of Level 2 SPF that was run.
Prefix or **Label**	Destination of the route.
L	IS-IS level: • **1**—Level 1 only • **2**—Level 2 only • **3**—Level 1 and Level 2
Version	Version of SPF that generated the route.
Metric	Metric value associated with the route.
Type	Metric type: **int** (internal) or **ext** (external).
Interface	Interface to the next hop.
Via	System identifier of the next hop, displayed as a name if possible.
ISO Routes	ISO routing table entries.
snpa	MAC address.

Sample Output

show isis route logical-system

```
user@host> show isis route logical-system ls1
IS-IS routing table             Current version: L1: 8 L2: 11
Prefix             L Version Metric Type Interface    Via
10.9.7.0/30        2       11     20 int  gr-0/2/0.0   h
10.9.201.1/32      2       11     60 int  gr-0/2/0.0   h
IPV6 Unicast IS-IS routing table             Current version: L1: 9 L2: 11
Prefix             L Version Metric Type Interface    Via
8009:3::a09:3200/126 2      11     20 int gr-0/2/0.0   h
```

show isis route floodgroup

```
user@R2> show isis route floodgroup 49.0001
IS-IS routing table             Current version: L1: 14 L2: 27
IPv4/IPv6 Routes
```


Copyright © 2018, Juniper Networks, Inc.

```
----------------
Prefix              L Version   Metric Type Interface       NH   Via
    Backup Score
0.0.0.0/0           1      14       10 int  ge-0/0/8.0      IPV4 R1
81.3.3.3/32         1      14       10 int  ge-0/0/8.0      IPV4 R1
128.220.17.202/32   1      14       10 int  ge-0/0/8.0      IPV4 R1
```

show isis route (CLNS)

```
user@host> show isis route
IS-IS routing table             Current version: L1: 10 L2: 8
IPv4/IPv6 Routes
Prefix              L Version   Metric Type Interface    Via
0.0.0.0/0           1      10       10 int  fe-0/0/1.0   ISIS.0
ISO Routes
Prefix  L   Version   Metric Type Interface    Via    snpa
0/0
        1        10       10 int  fe-0/0/1.0   isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0001/104
        1        10        0 int
47.0005.80ff.f800.0000.0108.0001.1921.6800.4001/152
        1        10       10 int  fe-0/0/1.0 isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0001.1921.6800.4002/152
        1        10       20 int  fe-0/0/1.0 isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0002/104
        1        10        0 int
47.0005.80ff.f800.0000.0108.0002.1921.6800.4001/152
        1        10       10 int  fe-0/0/1.0 isis.0 0:12:0:34:0:56
```

show isis route

```
user@host> show isis route

IS-IS routing table             Current version: L1: 4 L2: 13
IPv4/IPv6 Routes
----------------
Prefix                   L   Version  Metric Type Interface       NH   Via
10.255.71.52/32          2      13       10    int  ae0.0           IPV4 camaro

10.255.71.238/32         2      13       20    int  so-6/0/0.0      IPV4 olympic

                                                    as0.0           IPV4 glacier

10.255.71.239/32         2      13       20    int  so-6/0/0.0      IPV4 olympic

                                                    ae0.0           IPV4 camaro

10.255.71.242/32         2      13       10    int  as0.0           IPV4 glacier

10.255.71.243/32         2      13       10    int  so-6/0/0.0      IPV4 olympic

12.13.0.0/30             2      13       20    int  so-6/0/0.0      IPV4 olympic

12.15.0.0/30             2      13       20    int  so-6/0/0.0      IPV4 olympic

13.15.0.0/30             2      13       30    int  ae0.0           IPV4 camaro

                                                    so-6/0/0.0      IPV4 olympic

                                                    as0.0           IPV4 glacier
```

```
13.16.0.0/30            2      13       25   int  as0.0           IPV4 glacier

14.15.0.0/30            2      13       20   int  ae0.0           IPV4 camaro

192.2.1.0/30            2      13       30   int  so-6/0/0.0      IPV4 olympic

                                                  as0.0           IPV4 glacier

1eee::/64               2      13       30   int  so-6/0/0.0      IPV6 olympic

                                                  as0.0           IPV6 glacier

abcd::10:255:71:52/128   2      13       10   int  ae0.0           IPV6 camaro

abcd::10:255:71:238/128  2      13       20   int  so-6/0/0.0      IPV6 olympic

                                                  as0.0           IPV6 glacier

abcd::10:255:71:239/128  2      13       20   int  so-6/0/0.0      IPV6 olympic

                                                  ae0.0           IPV6 camaro

abcd::10:255:71:242/128  2      13       10   int  as0.0           IPV6 glacier

abcd::10:255:71:243/128  2      13       10   int  so-6/0/0.0      IPV6 olympic


MPLS Routes
-----------
Label              L Version   Metric Type Interface       NH     Via
300032 /52         2      38        0 int  lt-1/2/0.13     MPLS   Direct forward
 to 10.0.7.60(pro-bng3-c-E)
300048 /52         1      27        0 int  lt-1/2/0.12     MPLS   Direct forward
 to 10.0.6.60(pro-bng3-c-E)
300064 /52         1      27        0 int  lt-1/2/0.12     MPLS   Direct forward
 to 10.0.6.60(pro-bng3-c-E)
300080 /52         2      38        0 int  lt-1/2/0.12     MPLS   Direct forward
 to 10.0.6.60(pro-bng3-c-E)
300096 /52         2      38        0 int  lt-1/2/0.12     MPLS   Direct forward
 to 10.0.6.60(pro-bng3-c-E)
299920 /52         1      27        0 int  lt-1/2/0.14     MPLS   Direct forward
 to 10.0.10.70(pro-bng3-c-F)
299936 /52         1      27        0 int  lt-1/2/0.14     MPLS   Direct forward
 to 10.0.10.70(pro-bng3-c-F)
299952 /52         2      38        0 int  lt-1/2/0.14     MPLS   Direct forward
 to 10.0.10.70(pro-bng3-c-F)
299968 /52         2      38        0 int  lt-1/2/0.14     MPLS   Direct forward
 to 10.0.10.70(pro-bng3-c-F)
299984 /52         1      27        0 int  lt-1/2/0.13     MPLS   Direct forward
 to 10.0.7.60(pro-bng3-c-E)
300000 /52         1      27        0 int  lt-1/2/0.13     MPLS   Direct forward
 to 10.0.7.60(pro-bng3-c-E)
300016 /52         2      38        0 int  lt-1/2/0.13     MPLS   Direct forward
 to 10.0.7.60(pro-bng3-c-E)
```

 Copyright © 2018, Juniper Networks, Inc.

show isis spf

List of Syntax

Syntax on page 2027
Syntax (EX Series Switches) on page 2027

Syntax

show isis spf (brief | log | results)
<instance *instance-name*>
<level (1 | 2)>
<logical-system (all | *logical-system-name*)>
<topology (ipv4-multicast | ipv6-multicast | ipv6-unicast | unicast)>

Syntax (EX Series Switches)

show isis spf (brief | log | results)
<instance *instance-name*>
<level (1 | 2)>
<topology (ipv4-multicast | ipv6-multicast | ipv6-unicast | unicast)>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display information about IS-IS shortest-path-first (SPF) calculations.

Options

brief—Display an overview of SPF calculations.

log—Display the log of SPF calculations.

results—Display the results of SPF calculations.

instance *instance instance-name*—(Optional) Display SPF calculations for the specified routing instance.

level (1 | 2)—(Optional) Display SPF calculations for the specified IS-IS level.

log—Display the log of SPF calculations.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

results—Display the results of SPF calculations.

topology (ipv4-multicast | ipv6-multicast | ipv6-unicast | unicast)—(Optional) Display SPF calculations for the specified topology only.

Required Privilege Level

view

List of Sample Output

show isis spf log on page 2028
show isis spf results logical-system on page 2029
show isis spf results (CLNS) on page 2030

Output Fields

Table 162 on page 2028 describes the output fields for the **show isis spf** command. Output fields are listed in the approximate order in which they appear.

Table 162: show isis spf Output Fields

Field Name	Field Description
Node	System ID of a node.
Metric	Metric to the node.
Interface	Interface of the next hop.
Via	System ID of the next hop.
SNPA	Subnetwork point of attachment (MAC address of the next hop).
Start time	(**log** option only) Time that the SPF computation started.
Elapsed (secs)	(**log** option only) Length of time, in seconds, required to complete the SPF computation.
Count	(**log** option only) Number of times the SPF was triggered.
Reason	(**log** option only) Reason that the SPF computation was completed.

Sample Output

show isis spf log

```
user@host> show isis spf log logical-system lsl
 IS-IS level 1 SPF log:
Start time          Elapsed (secs) Count Reason
Fri Oct 31 12:41:18        0.000069    1 Reconfig
Fri Oct 31 12:41:18        0.000107    3 Updated LSP fix.00-00
Fri Oct 31 12:41:18        0.000050    3 Address change on so-1/2/2.0
Fri Oct 31 12:41:23        0.000033    1 Updated LSP fix.00-00
Fri Oct 31 12:41:28        0.000178    5 New adjacency scat on ge-1/1/0.0
Fri Oct 31 12:41:59        0.000060    1 Updated LSP fix.00-00
Fri Oct 31 12:42:30        0.000161    2 Multi area attachment change
Fri Oct 31 12:56:58        0.000198    1 Periodic SPF
Fri Oct 31 13:10:29        0.000209    1 Periodic SPF
 IS-IS level 2 SPF log:

Start time          Elapsed (secs) Count Reason
Fri Oct 31 12:41:18        0.000035    1 Reconfig
Fri Oct 31 12:41:18        0.000047    2 Updated LSP fix.00-00
Fri Oct 31 12:41:18        0.000043    5 Address change on gr-0/2/0.0
Fri Oct 31 12:41:23        0.000022    1 Updated LSP fix.00-00
Fri Oct 31 12:41:59        0.000144    3 New adjacency h on gr-0/2/0.0
Fri Oct 31 12:42:30        0.000257    3 New LSP skag.00-00
Fri Oct 31 12:54:37        0.000195    1 Periodic SPF
Fri Oct 31 12:55:50        0.000178    1 Updated LSP fix.00-00
Fri Oct 31 12:55:55        0.000174    1 Updated LSP h.00-00
Fri Oct 31 12:55:58        0.000176    1 Updated LSP skag.00-00
```


Copyright © 2018, Juniper Networks, Inc.

```
Fri Oct 31 13:08:14        0.000198    1 Periodic SPF
IPV6 Unicast IS-IS level 1 SPF log:

Start time          Elapsed (secs) Count Reason
Fri Oct 31 12:41:18        0.000028    1 Reconfig
Fri Oct 31 12:41:18        0.000043    3 Updated LSP fix.00-00
Fri Oct 31 12:41:18        0.000112    4 Updated LSP fix.00-00
Fri Oct 31 12:41:23        0.000059    1 Updated LSP fix.00-00
Fri Oct 31 12:41:25        0.000041    1 Updated LSP fix.00-00
Fri Oct 31 12:41:28        0.000103    5 New adjacency scat on ge-1/1/0.0
Fri Oct 31 12:41:59        0.000040    1 Updated LSP fix.00-00
Fri Oct 31 12:42:30        0.000118    2 Multi area attachment change
Fri Oct 31 12:56:08        0.000289    1 Periodic SPF
Fri Oct 31 13:11:07        0.000214    1 Periodic SPF
IPV6 Unicast IS-IS level 2 SPF log:

Start time          Elapsed (secs) Count Reason
Fri Oct 31 12:41:18        0.000027    1 Reconfig
Fri Oct 31 12:41:18        0.000039    2 Updated LSP fix.00-00
Fri Oct 31 12:41:18        0.000049    6 Updated LSP fix.00-00
Fri Oct 31 12:41:23        0.000025    1 Updated LSP fix.00-00
Fri Oct 31 12:41:25        0.000023    1 Updated LSP fix.00-00
Fri Oct 31 12:41:59        0.000087    3 New adjacency h on gr-0/2/0.0
Fri Oct 31 12:42:30        0.000123    3 New LSP skag.00-00
Fri Oct 31 12:55:50        0.000121    1 Updated LSP fix.00-00
Fri Oct 31 12:55:55        0.000121    1 Updated LSP h.00-00
Fri Oct 31 12:55:58        0.000121    1 Updated LSP skag.00-00
Fri Oct 31 13:09:46        0.000201    1 Periodic SPF
...
```

show isis spf results logical-system

```
user@host> show isis spf results logical-system ls1
IS-IS level 1 SPF results:
Node                Metric      Interface          Via              SNPA
scat.00             10          ge-1/1/0.0        scat             0:90:69:a6:48:9d
                    20          10.9.1.0/30
fix.02              10
fix.00              0
                    10          10.9.1.0/30
                    10          10.9.5.0/30
                    10          10.9.6.0/30
                    20          10.9.7.0/30
                    60          10.9.201.1/32
  3 nodes

 IS-IS level 2 SPF results:
Node                Metric      Interface          Via              SNPA
skag.00             20          gr-0/2/0.0        h
                    30          10.9.7.0/30
skag.02             20          gr-0/2/0.0        h
h.00                10          gr-0/2/0.0        h
                    20          10.9.6.0/30
                    20          10.9.7.0/30
                    60          10.9.201.1/32
fix.00              0
                    10          10.9.1.0/30
                    10          10.9.5.0/30
                    10          10.9.6.0/30
  4 nodes
```

```
IPV6 Unicast IS-IS level 1 SPF results:
Node             Metric     Interface        Via              SNPA
scat.00          10         ge-1/1/0.0       scat             0:90:69:a6:48:9d
                            ge-1/1/0.0       scat             0:90:69:a6:48:9d
                 20         8009:1::a09:1400/126
fix.02           10
fix.00           0
                 10         8009:1::a09:1400/126
                 10         8009:2::a09:1e00/126
                 20         8009:3::a09:3200/126
                 10         8009:4::a09:2800/126
  3 nodes

IPV6 Unicast IS-IS level 2 SPF results:
Node             Metric     Interface        Via              SNPA
skag.00          20         gr-0/2/0.0       h
                            gr-0/2/0.0       h
                 30         8009:3::a09:3200/126
skag.02          20         gr-0/2/0.0       h
                            gr-0/2/0.0       h
h.00             10         gr-0/2/0.0       h
                            gr-0/2/0.0       h
                 20         8009:3::a09:3200/126
                 20         8009:4::a09:2800/126
fix.00           0
                 10         8009:1::a09:1400/126
                 10         8009:2::a09:1e00/126
                 10         8009:4::a09:2800/126
  4 nodes

Multicast IS-IS level 1 SPF results:
Node             Metric     Interface        Via              SNPA
scat.00          10         ge-1/1/0.0       scat             0:90:69:a6:48:9d
fix.02           10
fix.00           0
  3 nodes

Multicast IS-IS level 2 SPF results:
Node             Metric     Interface        Via              SNPA
skag.00          20         gr-0/2/0.0       h
skag.02          20         gr-0/2/0.0       h
h.00             10         gr-0/2/0.0       h
fix.00           0
  4 nodes
...
```

show isis spf results (CLNS)

```
user@host> show isis spf results
IS-IS level 1 SPF results:
Node             Metric     Interface        Via              SNPA
skag.00 10                  fe-0/0/1.0       toothache        0:12:0:34:0:56
                            fe-0/0/1.0       toothache        0:12:0:34:0:56
                 20         192.168.37.64/29
                 10         1921.6800.4001
                 20         1921.6800.4002
pro1-a.02        10
pro1-a.00        0
                 0          10.255.245.1/32
                 10         192.168.37.64/29
                 0          1921.6800.4211
```


Copyright © 2018, Juniper Networks, Inc.

```
  3 nodes

IS-IS level 2 SPF results:
Node             Metric     Interface         Via             SNPA
skag.00 10                  fe-0/0/1.0       toothache       0:12:0:34:0:56
                            fe-0/0/1.0       toothache       0:12:0:34:0:56
                 20         10.255.245.1/32
                 20         192.168.37.64/29
                 20         47.0005.80ff.f800.0000.0109.0010/104
pro1-a.02        10
pro1-a.00        0
                 0          10.255.245.1/32
                 10         192.168.37.64/29
  3 nodes
```

show isis statistics

List of Syntax

Syntax on page 2032
Syntax (EX Series Switches and QFX Series) on page 2032

Syntax

show isis statistics
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches and QFX Series)

show isis statistics
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display statistics about IS-IS traffic.

Options

none—Display IS-IS traffic statistics for all routing instances.

instance *instance-name*—(Optional) Display statistics for the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- clear isis statistics on page 1978

List of Sample Output

show isis statistics on page 2034

Output Fields

Table 163 on page 2033 describes the output fields for the **show isis statistics** command. Output fields are listed in the approximate order in which they appear.

 Copyright © 2018, Juniper Networks, Inc.

Table 163: show isis statistics Output Fields

Field Name	Field Description
PDU type	PDU type: • **CSNP**—Complete sequence number PDUs contain a complete list of all link-state PDUs in the IS-IS database. CSNPs are sent periodically on all links, and the receiving systems use the information in the CSNP to update and synchronize their link-state PDU databases. The designated router multicasts CSNPs on broadcast links in place of sending explicit acknowledgments for each link-state PDU. • **IIH**—IS-IS hello packets are broadcast to discover the identity of neighboring IS-IS systems and to determine whether the neighbors are Level 1 or Level 2 intermediate systems. • **LSP**—Link-state PDUs contain information about the state of adjacencies to neighboring IS-IS systems. Link-state PDUs are flooded periodically throughout an area. • **PSNP**—Partial sequence number PDUs are sent multicast by a receiver when it detects that it is missing a link-state PDU (when its link-state PDU database is out of date). The receiver sends a PSNP to the system that transmitted the CSNP, effectively requesting that the missing link-state PDU be transmitted. That routing device, in turn, forwards the missing link-state PDU to the requesting routing device. • **Unknown**—The PDU type is unknown.
Received	Number of PDUs received since IS-IS started or since the statistics were set to zero.
Processed	Number of PDUs received less the number dropped.
Drops	Number of PDUs dropped.
Sent	Number of PDUs transmitted since IS-IS started or since the statistics were set to zero.
Rexmit	Number of PDUs retransmitted since IS-IS started or since the statistics were set to zero.
Total packets received/sent	Total number of PDUs received and transmitted since IS-IS started or since the statistics were set to zero.
SNP queue length	Number of CSPN and PSNP packets currently waiting in the queue for processing. This value is almost always **0**.
LSP queue length	Number of link-state PDUs waiting in the queue for processing. This value is almost always **0**.
SPF runs	Number of shortest-path-first (SPF) calculations that have been performed. If this number is incrementing rapidly, it indicates that the network is unstable.
Fragments rebuilt	Number of link-state PDU fragments that the local system has computed.
LSP regenerations	Number of link-state PDUs that have been regenerated. A link-state PDU is regenerated when it is nearing the end of its lifetime and it has not changed.
Purges initiated	Number of purges that the system initiated. A purge is initiated if the software decides that a link-state PDU must be removed from the network.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show isis statistics

```
user@host> show isis statistics
IS-IS statistics for merino:

PDU type      Received  Processed      Drops       Sent      Rexmit
LSP              12227      12227          0       8184         683
IIH             113808     113808          0     115817           0
CSNP            198868     198868          0     198934           0
PSNP              6985       6979          6       8274           0
Unknown              0          0          0          0           0
Totals          331888     331882          6     331209         683

Total packets received: 331888 Sent: 331892

SNP queue length:           0 Drops:          0
LSP queue length:           0 Drops:          0

SPF runs:                1014
Fragments rebuilt:       1038
LSP regenerations:        425
Purges initiated:           0
```

 Copyright © 2018, Juniper Networks, Inc.

CHAPTER 22

LLDP Operational Commands

clear lldp neighbors

Syntax

clear lldp neighbor
<interface *interface-name*>

Release Information

Command introduced in Junos OS Release 9.6.

Description

Clear information regarding all Link Layer Discovery Protocol (LLDP) neighbors or LLDP neighbors of the specified interface.

For information about interface names, see "Interface Naming Overview" on page 32. For information about interface names for TX Matrix routers, see *TX Matrix Router Chassis and Interface Names*. For information about FPC numbering on TX Matrix routers, see *Routing Matrix with a TX Matrix Router FPC Numbering*.

For information about interface names in the Junos Fusion technology, see *Understanding Junos Fusion Ports* .

Options

interface *interface-name*—(Optional) Clear the LLDP neighbors on the specified interface.

Required Privilege Level

clear

Related Documentation

- clear lldp statistics on page 2037

List of Sample Output

clear lldp neighbors on page 2036
clear lldp neighbors interface ge-0/1/1.0 on page 2036

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the **show lldp neighbors** command before and after clearing the LLDP neighbors to verify the clear operation.

Sample Output

clear lldp neighbors

```
user@switch> clear lldp neighbors
```

clear lldp neighbors interface ge-0/1/1.0

```
user@switch> clear lldp neighbors interface ge-0/1/1.0
```


Copyright © 2018, Juniper Networks, Inc.

clear lldp statistics

Syntax

clear lldpp neighbor
<interface *interface-name*>

Release Information

Command introduced in Junos OS Release 9.6.

Description

Clear all Link Layer Discovery Protocols (LLDP) statistics or LLDP statistics associated with the specified interface.

For information about interface names, see "Interface Naming Overview" on page 32. For information about interface names for TX Matrix routers, see *TX Matrix Router Chassis and Interface Names*. For information about FPC numbering on TX Matrix routers, see *Routing Matrix with a TX Matrix Router FPC Numbering*.

For information about interface names in the Junos Fusion technology, see *Understanding Junos Fusion Ports* .

Options

interface *interface-name*—(Optional) Clear LLDP statistics on the specified interface.

Required Privilege Level

clear

Related Documentation

- clear lldp neighbors on page 2036

List of Sample Output

clear lldp statistics on page 2037
clear lldp statistics interface ge-0/1/1.0 on page 2037

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the **show lldp statistics** command before and after clearing the LLDP statistics to verify the clear operation.

Sample Output

clear lldp statistics

```
user@switch> clear lldp statistics
```

clear lldp statistics interface ge-0/1/1.0

```
user@switch> clear lldp statistics interface ge-0/1/1.0
```

show lldp

Syntax

show lldp
<detail>

Release Information

Command introduced in Junos OS Release 9.6.

Description

Display information about the Link Layer Discovery Protocol (LLDP).

Options

detail—(Optional) Display the detailed output level.

Required Privilege Level

view

List of Sample Output

show lldp on page 2040
show lldp detail on page 2040

Output Fields

Table 164 on page 2038 describes the output fields for the **show lldp** command. Output fields are listed in the approximate order in which they appear.

Table 164: show lldp Output Fields

Field Name	Field Description
LLDP	Status of LLDP: **Enabled** or **Disabled**.
Advertisement interval	Value of the advertisement interval parameter.
Transmit delay	Value of the transmit delay parameter.
Hold timer	Value of the hold timer parameter.
Notification interval	Value of the notification interval parameter.
Config Trap Interval	Value of the configuration trap parameter.
Connection Hold timer	Value of the connection hold timer parameter.
Port ID TLV subtype	• *interface-name*—Indicates the interface name as the port information for the local device. • **locally-assigned**—Indicates that the sub-type for port ID TLV generation is locally assigned value of SNMP index of the interface. For more information about port ID TLV subtype, see *port-id-subtype*..

Copyright © 2018, Juniper Networks, Inc.

Table 164: show lldp Output Fields (continued)

Field Name	Field Description
Port Description TLV type	Following value used for port description TLV: • **interface-alias (ifAlias)**—Indicates that the *ifAlias* MIB object value is used to generate the port description TLV. • **interface-description (ifDescr)**—Indicates that the *ifDescr* MIB object value is used to generate the port description TLV. For more information about port description TLV type, see *port-description-type*.
Interface	Name of the interface for which LLDP configuration information is being reported For information about interface names, see "Interface Naming Overview" on page 32. For information about interface names for TX Matrix routers, see *TX Matrix Router Chassis and Interface Names*. For information about FPC numbering on TX Matrix routers, see *Routing Matrix with a TX Matrix Router FPC Numbering*.
Parent Interface	Name of the aggregated Ethernet interface, if any, to which the interface belongs.
LLDP	LLDP operating state. The state can be Enabled or Disabled.
LLDP-MED	LLDP-MED operating state. The state can be Enabled or Disabled.
Power Negotiation	LLDP power negotiation operating state. The state can be Enabled or Disabled.
LLDP basic TLVs supported	List of basic LLDP TLVs supported by this device (**detail** only).
LLDP 802 TLVs supported	List of IEEE 802.1 LLDP TLVs supported by this device (**detail** only).

Sample Output

show lldp

```
user@host> show lldp
LLDP                        : Enabled
Advertisement interval      : 30 seconds
Transmit delay              : 2 seconds
Hold timer                  : 120 seconds
Notification interval       : 0 Second(s)
Config Trap Interval        : 0 seconds
Connection Hold timer       : 300 seconds
Port ID TLV subtype         : locally-assigned
Port Description TLV type : interface-description (ifDescr)

Interface       Parent Interface    LLDP        LLDP-MED       Power Negotiation
all             -                   Enabled
```

Sample Output

show lldp detail

```
user@host> show lldp detail
LLDP                        : Enabled
Advertisement interval      : 30 seconds
Transmit delay              : 2 seconds
Hold timer                  : 120 seconds
Notification interval       : 0 Second(s)
Config Trap Interval        : 0 seconds
Connection Hold timer       : 300 seconds
Port ID TLV subtype         : locally-assigned
Port Description TLV type : interface-description (ifDescr)

Interface       Parent Interface    LLDP        LLDP-MED     Power Negotiation
  Neighbor count
all             -                   Enabled
  2

Interface       Parent Interface    Vlan-id     Vlan-name
xe-0/0/0        -                   4080        vlan-4080
xe-0/0/1        -                   4080        vlan-4080

Basic Management TLVs supported:
End Of LLDPDU, Chassis ID, Port ID, Time To Live, Port Description, System Name,
 System Description, System Capabilities, Management Address

Organizationally Specific TLVs supported:
Port VLAN tag, VLAN Name, MAC/PHY Configuration/Status, Link Aggregation,Maximum
 Frame Size
```


Copyright © 2018, Juniper Networks, Inc.

show lldp local-information

Syntax show lldp local-information

Release Information Command introduced in Junos OS Release 9.6.

Description Display local Link Layer Discovery Protocol (LLDP) information.

Options This command has no options.

Required Privilege Level view

List of Sample Output
show lldp local-information(Management Information Address Subtype is IPv4) on page 2043
show lldp local-information(Management Information Address Subtype is IPv6) on page 2043

Output Fields Table 165 on page 2041 describes the output fields for the **show lldp local-information** command. Output fields are listed in the approximate order in which they appear.

Table 165: show lldp local-information Output Fields

Field Name	Field Description
LLDP Local Information details	Information that follows pertains to the local system.
Chassis ID	List of chassis identifiers for local information.
System name	Local system name reported by LLDP.
System descr	Local system description reported by LLDP.
System Capabilities	Capabilities (such as **Bridge** or **Router**) that are **Supported** or **Enabled** by system on the interface.
Management Information	Listed by **Interface Name**, **Address Subtype** (such as **ipv4**, **ipv6**), **Address** (such as **192.168.168.229**, **1fd::1a10**), **Interface Number**, and **Interface Numbering Subtype**.
Interface Name	List of local interfaces. For information about interface names, see "Interface Naming Overview" on page 32. For information about interface names for TX Matrix routers, see *TX Matrix Router Chassis and Interface Names*. For information about FPC numbering on TX Matrix routers, see *Routing Matrix with a TX Matrix Router FPC Numbering*.
Parent Interface	Name of the **ae** interface to which the interface belongs
Interface ID	List of local interface identifiers.

Copyright © 2018, Juniper Networks, Inc.

Table 165: show lldp local-information Output Fields (continued)

Field Name	Field Description
Interface Description	List of local interface descriptions.
Status	List of interface conditions: **UP** or **DOWN**.

 Copyright © 2018, Juniper Networks, Inc.

Sample Output

show lldp local-information(Management Information Address Subtype is IPv4)

```
user@host> show lldp local-information
LLDP Local Information details

Chassis ID   : 64:87:88:65:37:c0
System name  : apg-hp1
System descr : Juniper Networks, Inc. mx240 , version 14.1I20131231_0701_builder
 [builder] Build date: 2013-12-31 07:13:42 UTC

System Capabilities
    Supported       : Bridge Router
    Enabled         : Bridge Router

Management Information
    Interface Name   : Unknown
    Address Subtype : IPv4(1)
    Address         : 10.216.97.103
    Interface Number               : 1
    Interface Numbering Subtype    : ifIndex(2)

Interface name   Parent Interface   Interface ID  Interface description   Status
fxp0                -                1             fxp0                     Up
me0                 -                33            me0                      Up
ge-2/0/0         ae0                1475          ge-2/0/0                 Up
ge-2/0/1         ae0                1476          ge-2/0/1                 Up
```

show lldp local-information(Management Information Address Subtype is IPv6)

```
user@host> show lldp local-information
LLDP Local Information details

Chassis ID   : ac:4b:c8:92:67:c0
System name  : apg-hp
System descr : Juniper Networks, Inc. mx240 , version 13.2-20131210.0 [builder]
               Build date: 2013-12-10 06:23:15 UTC

System Capabilities
    Supported       : Bridge Router
    Enabled         : Bridge Router

Management Information
    Interface Name   : fxp0
    Address Subtype : IPv6(2)
    Address         : 1fd::1a20
    Interface Number               : 1
    Interface Numbering Subtype    : ifIndex(2)

Interface name  Parent Interface   Interface ID   Interface description    Status
ge-1/2/4        -                  530            -                          Down
ge-1/2/5        -                  531            -                          Down
ge-1/2/2        -                  528            ge-1/2/2                    Up
ge-1/2/3        -                  529            ge-1/2/3                    Up
```

show lldp neighbors

Syntax

show lldp neighbors
<interface *interface-name*>

Release Information

Command introduced in Junos OS Release 9.6.

Description

Display information about LLDP neighbors.

For information about interface names, see "Interface Naming Overview" on page 32. For information about interface names for TX Matrix routers, see *TX Matrix Router Chassis and Interface Names*. For information about FPC numbering on TX Matrix routers, see *Routing Matrix with a TX Matrix Router FPC Numbering*.

For information about extended port names in the Junos Fusion technology, see *Understanding Junos Fusion Ports* .

Options

interface ***interface-name***—(Optional) Display the neighbor information about a particular physical interface.

NOTE: Starting with Junos OS Release 14.2, you can also display LLDP neighbor details for management interfaces, such as fxp or me, on MX Series routers.

Required Privilege Level

view

Related Documentation

- clear lldp neighbors on page 2036

List of Sample Output

show lldp neighbors on page 2047
show lldp neighbors interface ge-0/0/4 (Management Address is IPv4) on page 2047
show lldp neighbors interface ge-0/0/4 (Management Address is IPv6) on page 2048
show lldp neighbors (Management Ethernet Interfaces) on page 2049

Output Fields

Table 166 on page 2044 describes the output fields for the **show lldp neighbors** command. Output fields are listed in the approximate order in which they appear.

Table 166: show lldp neighbors Output Fields

Field Name	Field Description
LLDP Remote Devices Information	Information about remote devices.

Copyright © 2018, Juniper Networks, Inc.

Table 166: show lldp neighbors Output Fields (continued)

Field Name	Field Description
LocalInterface	List of local interfaces for which neighbor information is available.
ChassisId	List of chassis identifiers for neighbors.
PortInfo	List of port information gathered from neighbors. This could be the port identifier or port description.
SysName	List of system names gathered from neighbors.
LLDP Neighbor Information	Information about both local and neighbor systems on the interface (appears when the **interface** option is used).
Local Information	Information about local systems on the interface (appears when the **interface** option is used).
Neighbor Information	Information about both local and neighbor system on the interface (appears when the **interface** option is used).
Index	Local interface index (appears when the **interface** option is used).
Time Mark	Date and timestamp of information (appears when the **interface** option is used).
Time To Live	Number of seconds for which this information is valid (appears when the **interface** option is used).
Local Interface	Name of the local physical interface (appears when the **interface** option is used).
Parent Interface	Name of the **ae** interface to which the interface belongs
Local Port ID	Local port identifier (appears when the interface option is used).
Neighbor Information	Information about neighbor systems on the interface (appears when the **interface** option is used).
Chassis type	Type of chassis identifier supplied, such as **MAC address** (appears when the **interface** option is used).
Chassis ID	Chassis identifier of type listed (appears when the **interface** option is used).
Port type	Type of port identifier supplied, such as **local** (appears when the **interface** option is used).
Port ID	Port identifier of type listed (appears when the **interface** option is used).
Port description	Port description (appears when the **interface** option is used).

Copyright © 2018, Juniper Networks, Inc.

Table 166: show lldp neighbors Output Fields (continued)

Field Name	Field Description
System name	Name supplied by the system on the interface (appears when the **interface** option is used).
System Description	Description supplied by the system on the interface (appears when the **interface** option is used).
System Capabilities	Capabilities (such as **bridge** or **router**) that are **Supported** or **Enabled** by the system on the interface (appears when the **interface** option is used).
Management address	Details of the management address: **Address Type** (such as **ipv4** and **ipv6**), **Address** (such as **10.204.34.35**, **1fd::1a10**), **Interface Number**, **Interface Subtype**, and Organization Identifier (**OID**) (appears when the **interface** option is used).
Organization Info	One or more entries listing remote information by Organizationally Unique Identifier (**OUI**), **Subtype**, **Index**, and **Info** (appears when the **interface** option is used).

 Copyright © 2018, Juniper Networks, Inc.

Sample Output

show lldp neighbors

```
user@host> show lldp neighbors

Local Interface  Parent Interface  Chassis Id          Port info     System Name
ge-2/0/0         ae0               ac:4b:c8:92:67:c0   528           apg-hp
ge-2/0/1         ae0               ac:4b:c8:92:67:c0   529           apg-hp
```

Sample Output

show lldp neighbors interface ge-0/0/4 (Management Address is IPv4)

```
user@host> show lldp neighbors interface ge-0/0/4
LLDP Neighbor Information:
Local Information:
Index: 2 Time to live: 120 Time mark: Tue Dec 31 11:47:46 2013 Age: 15 secs
Local Interface    : ge-2/0/1
Parent Interface   : ae0
Local Port ID      : 1476
Ageout Count      : 0

Neighbour Information:
Chassis type       : Mac address
Chassis ID         : ac:4b:c8:92:67:c0
Port type          : Locally assigned
Port ID            : 529
Port description   : ge-1/2/3
System name        : apg-hp

System Description : Juniper Networks, Inc. mx240 , version 14.1-20131222.0
[builder] Build date: 2013-12-22 09:13:26 UTC


System capabilities
        Supported: Bridge Router
        Enabled  : Bridge Router

Management address
        Address Type      : IPv4(1)
        Address           : 10.216.98.57
        Interface Number  : 1
        Interface Subtype : ifIndex(2)
        OID               : 1.3.6.1.2.1.31.1.1.1.1.1.

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : MAC/PHY Configuration/Status (1)
       Info     : Autonegotiation [supported, enabled (0x3)], PMD Autonegotiation
 Capability (0x1d), MAU Type (0x0)
       Index    : 1

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : Link Aggregation (3)
       Info     : Aggregation Status (0x3), Aggregation Port ID (1694498816)
       Index    : 2

Organization Info
```

Copyright © 2018, Juniper Networks, Inc.

```
        OUI      : IEEE 802.3 Private (0x00120f)
        Subtype  : Maximum Frame Size (4)
        Info     : MTU Size (1518)
        Index    : 3
```

show lldp neighbors interface ge-0/0/4 (Management Address is IPv6)

```
user@host> show lldp neighbors interface ge-0/0/4
LLDP Neighbor Information:
Local Information:
Index: 1 Time to live: 120 Time mark: Thu Dec 12 07:19:45 2013 Age: 28 secs
Local Interface    : ge-1/2/2
Parent Interface   : -
Local Port ID      : 528
Ageout Count      : 0

Neighbour Information:
Chassis type       : Mac address
Chassis ID         : 64:87:88:65:37:c0
Port type          : Locally assigned
Port ID            : 1475
Port description   : ge-2/0/0
System name        : apg-hp1

System Description : Juniper Networks, Inc. mx240 , version 11.4R10 Build date:
2013-10-24 10:10:02 UTC


System capabilities
        Supported: Bridge Router
        Enabled  : Bridge Router

Management address
        Address Type      : IPv6(2)
        Address           : 1fd::1a10
        Interface Number  : 1
        Interface Subtype : ifIndex(2)
        OID               : 1.3.6.1.2.1.31.1.1.1.1.1.

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : MAC/PHY Configuration/Status (1)
      Info     : Autonegotiation [supported, enabled (0x3)], PMD Autonegotiation
 Capability (0x5), MAU Type (0x0)
       Index    : 1

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : Link Aggregation (3)
       Info     : Aggregation Status (0x1), Aggregation Port ID (0)
       Index    : 2

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : Maximum Frame Size (4)
       Info     : MTU Size (1518)
       Index    : 3

Organization Info
       OUI      : Ethernet Bridged (0x0080c2)
```

 Copyright © 2018, Juniper Networks, Inc.

```
Subtype   : VLAN Name (3)
Info      : VLAN ID (100), VLAN Name (vlan-100)
Index     : 4
```

show lldp neighbors (Management Ethernet Interfaces)

user@host> **show lldp neighbors**

```
Local Interface    Parent Interface    Chassis Id          Port info
System Name
fxp0               -                   78:fe:3d:ee:4e:00   151
x2-sw35
xe-0/0/0           -                   a8:d0:e5:50:26:c0   512
sitara
xe-0/0/1           -                   a8:d0:e5:50:26:c0   513
sitara
```

Copyright © 2018, Juniper Networks, Inc.

show lldp remote-global-statistics

Syntax show lldp remote-global-statistics

Release Information Command introduced in Junos OS Release 9.6.

Description Display remote Link Layer Discovery Protocol (LLDP) global statistics.

Options This command has no options.

Required Privilege Level view

List of Sample Output show lldp remote-global-statistics on page 2051

Output Fields Table 167 on page 2050 describes the output fields for the **show lldp remote-global-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 167: show lldp remote-global-statistics Output Fields

Field Name	Field Description
LLDP Remote Database Table Counters	Information about remote database table counters.
LastchangeTime	Time elapsed between LLDP agent startup and the last change to the remote database table information.
Inserts	Number of insertions made in the remote database table.
Deletes	Number of deletions made in the remote database table.
Drops	Number of LLDP frames dropped from the remote database table because of errors.
Ageouts	Number of remote database table entries that have aged out of the table.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show lldp remote-global-statistics

user@host> **show lldp remote-global-statistics**

```
user@host> show lldp remote-global-statistics
   LLDP Remote Database Table Counters
   LastchangeTime          Inserts     Deletes     Drops     Ageouts
   00:00:76 (76 sec)       192         0           0         0
```

show lldp statistics

Syntax

show lldp statistics
<interface *interface-name*>

Release Information

Command introduced in Junos OS Release 9.6.

Description

Display information about Link Layer Discovery Protocol (LLDP) statistics.

Options

interface *interface-name*—(Optional) Display the statistics about a particular physical interface.

NOTE: Starting with Junos OS Release 14.2, you can also display LLDP statistical details for management interfaces, such as fxp or me, on MX Series routers.

Required Privilege Level

view

Related Documentation

- clear lldp statistics on page 2037

List of Sample Output

show lldp statistics on page 2054
show lldp statistics interface ge-0/1/1 on page 2054

Output Fields

Table 168 on page 2052 describes the output fields for the **show lldp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 168: show lldp statistics Output Fields

Field Name	Field Description
Interface	Interface name. For information about interface names, see "Interface Naming Overview" on page 32. For information about interface names for TX Matrix routers, see *TX Matrix Router Chassis and Interface Names*. For information about FPC numbering on TX Matrix routers, see *Routing Matrix with a TX Matrix Router FPC Numbering*. For information about extended port names in the Junos Fusion technology, see *Understanding Junos Fusion Ports* .
Received	Number of LLDP frames received on this interface.
Transmitted	Number of LLDP frames sent on this interface.

Copyright © 2018, Juniper Networks, Inc.

Table 168: show lldp statistics Output Fields (continued)

Field Name	Field Description
Unknown-TLVs	Number of LLDP frames with unsupported content received on this interface.
With-Errors	Number of LLDP frames with errors received on this interface.
Discarded	Number of LLDP frames received on this interface that were discarded because of problems.
Transmitted	Total number of LLDP frames that were transmitted on an interface.
Untransmitted	Total number of LLDP frames that were untransmitted on an interface.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show lldp statistics

```
user@host> show lldp statistics

Interface     Parent Interface  Received  Unknown TLVs  With Errors
xe-3/0/0.0    ae31.0            1564      0             0
xe-3/0/1.0    ae31.0            1564      0             0
xe-3/0/2.0    ae31.0            1565      0             0
xe-3/0/3.0    ae31.0            1566      0             0
xe-3/0/4.0    ae31.0            1598      0             0
xe-3/0/5.0    ae31.0            1598      0             0
xe-3/0/6.0    ae31.0            1596      0             0
xe-3/0/7.0    ae31.0            1597      0             0
xe-5/0/6.0    -                 0         0             0
xe-5/0/7.0    -                 0         0             0


Discarded TLVs  Transmitted  Untransmitted
0               3044         1
0               3044         1
0               3044         1
0               3044         1
0               3075         1
0               3075         1
0               3075         1
0               3075         1
0               17312        0
0               17312        0
```

Sample Output

show lldp statistics interface ge-0/1/1

```
user@host> show lldp statistics interface ge-0/1/1
Interface Received Transmitted Unknown-TLVs With-Errors Discarded
--------- -------- ----------- ------------ ----------- ---------
ge-0/1/1  544      540         0            0           0
```


Copyright © 2018, Juniper Networks, Inc.

CHAPTER 23

MVRP Operational Commands

show mvrp

Syntax show mvrp

Release Information Command introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 10.1 for MX Series routers.
Command introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
Command introduced in Junos OS Release 15.1X49-D70 for SRX Series devices.

Description Display Multiple VLAN Registration Protocol (MVRP) configuration information.

Required Privilege Level view

Related Documentation

- *Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches*
- *Verifying That MVRP Is Working Correctly on Switches*
- show mvrp statistics on page 2067
- show mvrp applicant-state on page 2059
- show mvrp dynamic-vlan-memberships on page 2061
- show mvrp interface on page 2063
- show mvrp registration-state on page 2065
- show mvrp statistics on page 2067
- show mvrp applicant-state on page 2059
- show mvrp dynamic-vlan-memberships on page 2061
- show mvrp interface on page 2063
- show mvrp registration-state on page 2065

List of Sample Output show mvrp (EX Series switches and MX Series routers) on page 2057
show mvrp (SRX1500, SRX300, SRX550M, SRX345, SRX340 and SRX320) on page 2057
show mvrp (EX Series switches) on page 2058

Output Fields Table 169 on page 2056 lists the output fields for the **show mvrp** command. Output fields are listed in the approximate order in which they appear.

Table 169: show mvrp Output Fields

Field Name	Field Description
MVRP dynamic VLAN creation	Displays whether global MVRP dynamic VLAN creation is **Enabled** or **Disabled**.

 Copyright © 2018, Juniper Networks, Inc.

Table 169: show mvrp Output Fields (continued)

Field Name	Field Description
Global MVRP configuration	Displays global MVRP information: • **MVRP status**—Displays whether MVRP is **Enabled** or **Disabled**. • **MVRP dynamic vlan creation**—Displays whether global MVRP dynamic VLAN creation is **Dnabled** or **Disabled**.
MVRP BPDU MAC address	Displays the multicast media access control (MAC) address for MVRP. If configured, the provider MVRP multicast MAC address is used; otherwise, the customer MVRP multicast MAC address is used.
MVRP timers (ms)	Displays MVRP timer information: • **Interface**—The interface on which MVRP is configured. • **Join**—The maximum number of milliseconds the interfaces must wait before sending VLAN advertisements. • **Leave**—The number of milliseconds an interface must wait after receiving a Leave message to remove the interface from the VLAN specified in the message. • **LeaveAll**— The interval at which LeaveAll messages are sent on interfaces. LeaveAll messages maintain current MVRP VLAN membership information in the network.
Interface based configuration	Displays interface-specific MVRP information: • **Interface**—The interface on which MVRP is configured. • **Status**—Displays whether MVRP is **Enabled** or **Disabled**. • **Registration**—Displays whether registration for the interface is **Forbidden** or **Normal**. • **Dynamic VLAN Creation**—Displays whether interface dynamic VLAN creation is **Enabled** or **Disabled**.

Sample Output

show mvrp (EX Series switches and MX Series routers)

```
user@host> show mvrp
MVRP configuration for routing instance 'default-switch'
MVRP dynamic VLAN creation : Enabled
MVRP BPDU MAC address      : Customer bridge group (01-80-C2-00-00-21)
MVRP timers (ms)
  Interface       Join   Leave  LeaveAll
  ge-11/2/8        200     800     10000
  ge-11/0/9        200     800     10000
  ge-11/3/0        200     800     10000
```

Sample Output

show mvrp (SRX1500, SRX300, SRX550M, SRX345, SRX340 and SRX320)

```
user@host> show mvrp
MVRP configuration for routing instance 'default-switch'
MVRP dynamic VLAN creation : Enabled
MVRP BPDU MAC address      : Customer bridge group (00-00-5E-00-53-00)
MVRP timers (ms)
  Interface       Join   Leave  LeaveAll
  ge-0/0/1        200     800     60
```

Sample Output

show mvrp (EX Series switches)

```
user@switch> show mvrp

Global MVRP configuration
  MVRP status               : Enabled
  MVRP dynamic vlan creation: Enabled
  MVRP Timers (ms):
   Interface        Join   Leave   LeaveAll
   --------------   ----   -----   --------
   all              200    600     10000
   xe-0/1/1.0       200    600     10000

  Interface based configuration:
   Interface        Status     Registration   Dynamic VLAN Creation
   --------------   --------   ------------   ---------------------
   all       Disabled   Normal     Enabled
   xe-0/1/1.0     Enabled    Normal         Enabled
```

 Copyright © 2018, Juniper Networks, Inc.

show mvrp applicant-state

Syntax show mvrp applicant-state

Release Information Command introduced in Junos OS Release 10.1.
Command introduced in Junos OS Release 15.1X49-D70 for SRX Series devices.

Description For MX Series routers, EX Series switches, SRX1500, SRX300, SRX550M, SRX345, SRX340, and SRX320, display Multiple VLAN Registration Protocol (MVRP) applicant state information.

Required Privilege Level view

Related Documentation

- show mvrp on page 2056
- show mvrp interface on page 2063
- show mvrp registration-state on page 2065
- show mvrp statistics on page 2067
- show mvrp interface on page 2063
- show mvrp registration-state on page 2065

List of Sample Output

show mvrp applicant-state (EX Series and MX Series) on page 2060
show mvrp applicant-state on page 2060

Output Fields Table 170 on page 2059 lists the output fields for the **show mvrp applicant-state** command. Output fields are listed in the approximate order in which they appear.

Table 170: show mvrp applicant-state Output Fields

Field Name	Field Description
VLAN Id	Displays the VLAN ID number.
Interface	Displays the interface number associated with the VLAN ID.

Copyright © 2018, Juniper Networks, Inc.

Table 170: show mvrp applicant-state Output Fields (continued)

Field Name	Field Description
State	Displays one of the following MVRP registrar states: • VO— Very anxious observer. • VP —Very anxious passive. • VA —Very anxious new. • AN —Anxious new. • AA —Anxious active. • QA —Quiet active. • LA —Leaving active. • AO —Anxious observer. • QO —Quiet observer. • LO —Leaving observer. • AP —Anxious passive. • QA —Quiet passive.

Sample Output (EX Series and MX Series)

show mvrp applicant-state (EX Series and MX Series)

```
user@host> show mvrp applicant-state
MVRP applicant state for routing instance 'default-switch'
 (VO) Very anxious observer, (VP) Very anxious passive, (VA) Very anxious new,
 (AN) Anxious new, (AA) Anxious active, (QA) Quiet active, (LA) Leaving active,
 (AO) Anxious observer, (QO) Quiet observer, (LO) Leaving observer,
 (AP) Anxious passive, (QP) Quiet passive

VLAN Id      Interface          State
    100      ge-11/3/0          Declaring (QA)
    200      ge-11/3/0          Declaring (QA)
    300      ge-11/3/0          Declaring (QA)
```

Sample Output (SRX1500, SRX300, SRX550M, SRX345, SRX340, and SRX320)

show mvrp applicant-state

```
user@host> show mvrp applicant-state
MVRP applicant state for routing instance 'default-switch'
 (VO) Very anxious observer, (VP) Very anxious passive, (VA) Very anxious new,
 (AN) Anxious new, (AA) Anxious active, (QA) Quiet active, (LA) Leaving active,
 (AO) Anxious observer, (QO) Quiet observer, (LO) Leaving observer,
 (AP) Anxious passive, (QP) Quiet passive

VLAN Id      Interface          State
    1        ge-0/0/1          Idle (VO)
    30       ge-0/0/1          Idle (VO)
    40       ge-0/0/1          Idle (VO)
    50       ge-0/0/1          Idle (VO)
    100      ge-0/0/1          Idle (VO)
```


Copyright © 2018, Juniper Networks, Inc.

show mvrp dynamic-vlan-memberships

Syntax show mvrp dynamic-vlan-memberships

Release Information Command introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 10.1 for MX Series routers.
Command introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
Command introduced in Junos OS Release 15.1X49-D70 for SRX Series devices.

Description Display all VLANs that have been created dynamically using Multiple VLAN Registration Protocol (MVRP) on the router, switch, or SRX Series device.

Required Privilege Level clear

Related Documentation

- *Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches*
- *Verifying That MVRP Is Working Correctly on Switches*
- show mvrp on page 2056
- show mvrp applicant-state on page 2059
- show mvrp interface on page 2063
- show mvrp registration-state on page 2065
- show mvrp registration-state on page 2065
- show mvrp statistics on page 2067

List of Sample Output
show mvrp dynamic-vlan-memberships (MX Series and EX Series) on page 2061
show mvrp dynamic-vlan-memberships (EX Series) on page 2062
show mvrp dynamic-vlan-memberships on page 2062

Output Fields Table 171 on page 2061 lists the output fields for the **show mvrp dynamic-vlan-memberships** command on MX Series routers and EX Series switches. Output fields are listed in the approximate order in which they appear.

Table 171: show mvrp dynamic-vlan-memberships Output Fields

Field Name	Field Description
VLAN Id	The VLAN ID of the dynamically created VLAN.
Interfaces	The interface or interfaces that are bound to the dynamically created VLAN.

Sample Output (MX Series Routers and EX Series Switches)

show mvrp dynamic-vlan-

memberships (MX Series and EX Series)

```
user@host> show mvrp dynamic-vlan-memberships
MVRP dynamic vlans for routing instance 'default-switch'
 (s) static vlan, (f) fixed registration

VLAN Id            Interfaces
    100 (s)        ge-11/3/0
    200 (s)        ge-11/3/0
    300 (s)
```

Sample Output (EX Series Switches)

show mvrp dynamic-vlan-memberships (EX Series)

```
user@switch> show mvrp dynamic-vlan-memberships
VLAN Name             Interfaces
-------------------       ----------------
__mvrp_100__              xe-0/1/1.0
                          xe-0/1/0.0
__mvrp_200__              xe-0/1/1.0
                          xe-0/1/0.0
__mvrp_300__              xe-0/1/1.0
```

Sample Output (SRX1500, SRX300, SRX550M, SRX345, SRX340, SRX320)

show mvrp dynamic-vlan-memberships

```
user@host> show mvrp dynamic-vlan-memberships
MVRP dynamic vlans for routing instance 'default-switch'
 (s) static vlan, (f) fixed registration

VLAN Id            Interfaces
    1  (s)
    30 (s)
    40 (s)        ge-0/0/1
    50 (s)        ge-0/0/1
   100 (s)        ge-0/0/1 (f)
```

 Copyright © 2018, Juniper Networks, Inc.

show mvrp interface

Syntax show mvrp interface

Release Information Command introduced in Junos OS Release 10.1.
Command introduced in Junos OS Release 15.1X49-D70 for SRX Series devices.

Description Display Multiple VLAN Registration Protocol (MVRP) interface-specific information.

Required Privilege Level view

Related Documentation

- show mvrp on page 2056
- show mvrp applicant-state on page 2059
- show mvrp dynamic-vlan-memberships on page 2061
- show mvrp registration-state on page 2065
- show mvrp registration-state on page 2065
- show mvrp statistics on page 2067

List of Sample Output
show mvrp interface on page 2063
show mvrp interface on page 2064

Output Fields Table 172 on page 2063 lists the output fields for the **show mvrp interface** command. Output fields are listed in the approximate order in which they appear.

Table 172: show mvrp interface Output Fields

Field Name	Field Description
Interface	Interface on which MVRP is configured.
Status	Status of the MVRP: **Enabled** or **Disabled**.
Registration Mode	Registration for the interface: **Fixed**, **Forbidden**, or **Normal**.
Applicant Mode	Applicant mode.

Sample Output (MX Series Routers and SX Series Switches)

show mvrp interface

```
user@host> show mvrp interface
MVRP interface information for routing instance 'default-switch'

Interface        Status        Registration        Applicant
                               Mode                Mode
```

Copyright © 2018, Juniper Networks, Inc.

```
ge-11/2/8        Enabled      Normal           Normal
ge-11/0/9        Enabled      Normal           Normal
ge-11/3/0        Enabled      Normal           Normal
```

Sample Output (SRX1500, SRX300, SRX550M, SRX345, SRX340 and SRX320)

show mvrp interface

```
user@host> show mvrp interface
MVRP interface information for routing instance 'default-switch'

Interface        Status       Registration     Applicant
                 Mode         Mode             Mode
ge-0/0/1         Enabled      Normal           Normal
```

 Copyright © 2018, Juniper Networks, Inc.

show mvrp registration-state

Syntax show mvrp registration-state

Release Information Command introduced in Junos OS Release 10.1.
Command introduced in Junos OS Release 15.1X49-D70 for SRX Series devices.

Description For MX Series routers, EX Series switches and SRX1500, SRX300, SRX550M, SRX345, SRX340 and SRX320, display Multiple VLAN Registration Protocol (MVRP) registration state information.

Required Privilege Level view

Related Documentation

- show mvrp on page 2056
- show mvrp dynamic-vlan-memberships on page 2061
- show mvrp interface on page 2063
- show mvrp statistics on page 2067

List of Sample Output show mvrp registration-state (EX Series and MX Series) on page 2066
show mvrp registration-state (SRX1500, SRX300, SRX550M, SRX345, SRX340 and SRX320) on page 2066

Output Fields Table 173 on page 2065 lists the output fields for the **show mvrp registration-state** command. Output fields are listed in the approximate order in which they appear.

Table 173: show mvrp registration-state Output Fields

Field Name	Field Description
VLAN Id	Displays the VLAN ID number.
Interface	Displays the interface number associated with the VLAN ID.
Registrar State	Displays whether the registrar state is Registered or Empty.
Forced State	Displays whether the forced state is Registered or Empty.
Managed State	Displays one of the following states: • **fixed**—VLANs always stay in a registered state and are declared as such on all other forwarding ports. • **normal** —VLANs participate in the MVRP protocol and honor incoming join requests normally. • **forbidden** —VLANs ignore the incoming join requests and always stay in an unregistered state.
STP State	Displays whether the Spanning Tree Protocol (STP) is Blocking or Forwarding.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show mvrp registration-state (EX Series and MX Series)

```
user@host> show mvrp registration-state
MVRP registration state for routing instance 'default-switch'

VLAN Id   Interface   Registrar   Forced      Managed    STP
                      State       State       State      State
    100   ge-11/2/8   Empty       Registered  Fixed      Forwarding
          ge-11/0/9   Empty       Empty       Normal     Forwarding
          ge-11/3/0   Registered  Registered  Normal     Forwarding
    101   ge-11/2/8   Empty       Registered  Fixed      Forwarding
          ge-11/0/9   Empty       Empty       Normal     Forwarding
          ge-11/3/0   Registered  Registered  Normal     Forwarding
```

Sample Output

show mvrp registration-state (SRX1500, SRX300, SRX550M, SRX345, SRX340 and SRX320)

```
user@host> show mvrp registration-state
MVRP registration state for routing instance 'default-switch'

VLAN Id   Interface   Registrar   Forced      Managed    STP
                      State       State       State      State
    1     ge-0/0/1    Empty       Empty       Normal     Forwarding
    30    ge-0/0/1    Empty       Empty       Normal     Forwarding
    40    ge-0/0/1    Registered  Registered  Normal     Forwarding
    50    ge-0/0/1    Registered  Registered  Normal     Forwarding
   100    ge-0/0/1    Empty       Registered  Fixed      Forwarding
```


Copyright © 2018, Juniper Networks, Inc.

show mvrp statistics

List of Syntax

Syntax (EX Series Switches) on page 2067
Syntax (Switches with ELS Support) on page 2067
Syntax (SRX Devices) on page 2067

Syntax (EX Series Switches)

show mvrp statistics
<interface *interface-name*>

Syntax (Switches with ELS Support)

show mvrp statistics
<interface *interface-name*>
<routing-instance *routing-instance-name*>

Syntax (SRX Devices)

show mvrp statistics

Release Information

Command introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 13.2X50-D10 (ELS).
Command introduced in Junos OS Release 15.1X49-D70 for SRX Series devices.

Description

Display Multiple VLAN Registration Protocol (MVRP) statistics in the form of Multiple Registration Protocol data unit (MRPDU) messages.

Options

none—Show MVRP statistics for all interfaces on the switch.

interface *interface-name*—(Optional) Show MVRP statistics for the specified interface.

Required Privilege Level

view

Related Documentation

- show mvrp on page 2056
- *clear mvrp statistics*
- *Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches*
- *Verifying That MVRP Is Working Correctly on Switches*
- *Verifying That MVRP Is Working Correctly on EX Series Switches with ELS Support*

List of Sample Output

show mvrp statistics interface xe-0/1/1.0 on page 2070
show mvrp statistics on page 2070
show mvrp statistics (SRX Devices) on page 2070

Output Fields

Table 171 on page 2061 lists the output fields for the **show mvrp statistics** command on EX Series switches. Output fields are listed in the approximate order in which they appear.

Copyright © 2018, Juniper Networks, Inc.

Table 174: show mvrp statistics Output Fields

Field Name	Field Description
MRPDU received	Number of MRPDU messages received on the switch.
Invalid PDU received	Number of invalid MRPDU messages received on the switch.
New received	Number of new messages received on the switch.
Join Empty received	Number of MRP JoinEmpty messages received on the switch. Either this value or the value for *JoinIn received* should increase when the value for *MRPDU received* increases. If this value is not incrementing when it should, you might have a Junos OS release version compatibility issue. To fix a version compatibility issue, see *Configuring Multiple VLAN Registration Protocol (MVRP) on Switches* .
Join In received	Number of MRP JoinIn messages received on the switch. Either this value or the value for *JoinEmpty received* should increase when the value for *MRPDU received* increases. If this value is not incrementing when it should, you might have a Junos OS release version compatibility issue. To fix a version compatibility issue, see *Configuring Multiple VLAN Registration Protocol (MVRP) on Switches* .
Empty received	Number of MRP Empty messages received on the switch.
In received	Number of MRP In messages received on the switch.
Leave received	Number of MRP Leave messages received on the switch.
LeaveAll received	Number of LeaveAll messages received on the switch.
MRPDU transmitted	Number of MRPDU messages transmitted from the switch.
MRPDU transmit failures	Number of MRPDU transmit failures from the switch.
New transmitted	Number of new messages transmitted from the switch.
Join Empty transmitted	Number of JoinEmpty messages sent from the switch.
Join In transmitted	Number of MRP JoinIn messages sent from the switch.
Empty transmitted	Number of MRP Empty messages sent from the switch.
In transmitted	Number of MRP In messages sent from the switch.
Leave transmitted	Number of MRP Leave Empty messages sent from the switch.
LeaveAll transmitted	Number of MRP LeaveAll messages sent from the switch.

Table 175 on page 2069 lists the output fields for the **show mvrp statistics** command on SRX devices. Output fields are listed in the approximate order in which they appear.

 Copyright © 2018, Juniper Networks, Inc.

Table 175: show mvrp statistics Output Fields

Field Name	Field Description
Interface name	Interface for which MVRP statistics are displayed.
VLAN IDs registered	Number of Virtual LAN (VLAN) IDs registered.
Sent MVRP PDUs	Number of MRPDU messages transmitted from the switch.
Received MVRP PDUs without error	Number of MRPDU messages received on the switch.
Received MVRP PDUs with error	Number of invalid MRPDU messages received on the switch.
Transmitted Join Emply	Number of JoinEmpty messages sent from the switch.
Transmitted Leave All	Number of MRP LeaveAll messages sent from the switch.
Received Join In	Number of MRP JoinIn messages received on the switch. Either this value or the value for **Received Join Empty** should increase when the value for **Received MVRP PDUs without error** increases. If this value is not incrementing when it should, you might have a Junos OS release compatibility issue. To resolve the issue, see *Configuring Multiple VLAN Registration Protocol (MVRP) on Switches* .
Transmitted Join In	Number of MRP JoinIn messages sent from the switch.
Transmitted Empty	Number of MRP Empty messages sent from the switch.
Transmitted Leave	Number of MRP LeaveEmpty messages sent from the switch.
Transmitted In	Number of MRP In messages sent from the switch.
Transmitted New	Number of New messages transmitted from the switch.
Received Leave All	Number of LeaveAll messages received on the switch.
Received Leave	Number of MRP Leave messages received on the switch.
Received In	Number of MRP In messages received on the switch.
Received Empty	Number of MRP Empty messages received on the switch.
Received Join Empty	Number of MRP JoinEmpty messages received on the switch. Either this value or the value for **Received Join In** should increase when the value for **Received MVRP PDUs without error** increases. If this value is not incrementing when it should, you might have a Junos OS release compatibility issue. To resolve the issue, see *Configuring Multiple VLAN Registration Protocol (MVRP) on Switches* .
Received New	Number of New messages received on the switch.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show mvrp statistics interface xe-0/1/1.0

```
user@switch> show mvrp statistics interface xe-0/1/1.0
MVRP statistics
  MRPDU received                : 3342
  Invalid PDU received          : 0
  New received                  : 2
  Join Empty received           : 1116
  Join In received              : 2219
  Empty received                : 2
  In received                   : 2
  Leave received                : 1
  LeaveAll received             : 1117
  MRPDU transmitted             : 3280
  MRPDU transmit failures       : 0
  New transmitted               : 0
  Join Empty transmitted        : 1114
  Join In transmitted           : 2163
  Empty transmitted             : 1
  In transmitted                : 1
  Leave transmitted             : 1
  LeaveAll transmitted          : 1111
```

show mvrp statistics

```
user@host> show mvrp statistics
MVRP statistics for routing instance 'default-switch'

Interface name                  : xe-0/1/1
VLAN IDs registered             : 117
Sent MVRP PDUs                  : 118824
Received MVRP PDUs without error: 118848
Received MVRP PDUs with error   : 0
Transmitted Join Empty          : 5229
Transmitted Leave All           : 2
Recieved Join In                : 11884924
Transmitted Join In             : 1835
Transmitted Empty               : 93606408
Transmitted Leave               : 888
Transmitted In                  : 13780024
Transmitted New                 : 2692
Received Leave All              : 118761
Received Leave                  : 97
Received In                     : 3869
Received Empty                  : 828
Received Join Empty             : 2020152
Received New                    : 224
...
```

show mvrp statistics (SRX Devices)

```
user@host> show mvrp statistics
MVRP statistics for routing instance 'default-switch'

Interface name                  : ge-0/0/1
VLAN IDs registered             : 2
```

 Copyright © 2018, Juniper Networks, Inc.

```
Sent MVRP PDUs                  : 41
Received MVRP PDUs without error: 28
Received MVRP PDUs with error   : 0
Transmitted Join Empty          : 0
Transmitted Leave All           : 20
Received Join In                : 0
Transmitted Join In             : 0
Transmitted Empty               : 114
Transmitted Leave               : 0
Transmitted In                  : 10
Transmitted New                 : 0
Received Leave All              : 1
Received Leave                  : 0
Received In                     : 0
Received Empty                  : 67
Received Join Empty             : 24
Received New                    : 0
```

 Copyright © 2018, Juniper Networks, Inc.

CHAPTER 24

OSPF Operational Commands

clear (ospf | ospf3) database

List of Syntax

Syntax on page 2074
Syntax (EX Series Switch and QFX Series) on page 2074

Syntax

clear (ospf | ospf3) database
<all>
<advertising-router (*router-id* | self)>
<area *area-id*>
<asbrsummary>
<external>
<instance *instance-name*>
<inter-area-prefix>
<inter-area-router>
<intra-area-prefix>
<link-local>
<logical-system (all | *logical-system-name*)>
<lsa-id *lsa-id*>
<netsummary>
<network>
<nssa>
<opaque-area>
<purge>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>
<router>

Syntax (EX Series Switch and QFX Series)

clear (ospf | ospf3) database
<all>
<advertising-router (*router-id* | self)>
<area *area-id*>
<asbrsummary>
<external>
<instance *instance-name*>
<inter-area-prefix>
<inter-area-router>
<intra-area-prefix>
<link-local>
<lsa-id *lsa-id*>
<netsummary>
<network>
<nssa>
<opaque-area>
<purge>
<router>

Release Information

Command introduced before Junos OS Release 7.4.
advertising-router ***router-id***, **netsummary**, **network**, **nssa**, **opaque-area**, and **router** options added in Junos OS Release 8.3. You must use the **purge** command with these options.
area ***area-id*** option added in Junos OS Release 8.3.
Command introduced in Junos OS Release 9.0 for EX Series switches.
realm option added in Junos OS Release 9.2.
advertising-router (***router-id*** | **self**) option added in Junos OS Release 9.5.

Copyright © 2018, Juniper Networks, Inc.

advertising-router (*router-id* | **self**) option introduced in Junos OS Release 9.5 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
purge option (and all options that are dependent on the **purge** option) hidden in Junos OS Release 13.3.

Description With the master Routing Engine, delete entries in the Open Shortest Path First (OSPF) link-state advertisement (LSA) database. With the backup Routing Engine, delete the OSPF LSA database and sync the new database with the master Routing Engine.

CAUTION: You can also use the purge command with any of the options to discard rather than delete the specified LSA entries. This command is useful only for testing. Use it with care, because it causes significant network disruption.

Options **all**—Delete all LSAs other than the system's own LSAs, which are regenerated. To resynchronize the database, the system destroys all adjacent neighbors that are in the state **EXSTART** or higher. The neighbors are then reacquired and the databases are synchronized.

advertising-router (*router-id* | self)—(Optional) Discard entries for the LSA entries advertised by the specified routing device or by this routing device.

area *area-id*—(Optional) Discard entries for the LSAs in the specified area.

asbrsummary—(Optional) Discard summary AS boundary router LSA entries.

external—(Optional) Discard external LSAs.

instance *instance-name*—(Optional) Delete or discard entries for the specified routing instance only.

inter-area-prefix—(OSPFv3 only) (Optional) Discard interarea prefix LSAs.

inter-area-router—(OSPFv3 only) (Optional) Discard interarea router LSAs.

intra-area-prefix—(OSPFv3 only) (Optional) Discard intra-area prefix LSAs.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

link-local—(Optional) Delete link-local LSAs.

lsa-id *lsa-id*—(Optional) Discard the LSA entries with the specified LSA identifier.

netsummary—(Optional) Discard summary network LSAs.

network—(Optional) Discard network LSAs.

Copyright © 2018, Juniper Networks, Inc.

nssa—(Optional) Discard not-so-stubby area (NSSA) LSAs.

opaque-area—(Optional) Discard opaque area-scope LSAs.

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)—(OSPFv3 only) (Optional) Delete the entries for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

router—(Optional) Discard router LSAs.

purge—(Optional) Discard all entries in the link-state advertisement database. All link-state advertisements are set to **MAXAGE** and are flooded. The database is repopulated when the originators of the link-state advertisements receive the **MAXAGE** link-state advertisements and reissue them.

Required Privilege Level clear

Related Documentation

- show ospf database on page 2100
- show ospf3 database on page 2111

List of Sample Output clear ospf database all on page 2076

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ospf database all

```
user@host> clear ospf database all
```

 Copyright © 2018, Juniper Networks, Inc.

clear (ospf | ospf3) io-statistics

List of Syntax

Syntax on page 2077
Syntax (EX Series Switch and QFX Series) on page 2077

Syntax

clear (ospf | osfp3) io-statistics
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and QFX Series)

clear (ospf | osfp3) io-statistics

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Clear Open Shortest Path First (OSPF) input and output statistics.

Options

none—Clear OSPF input and output statistics.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

List of Sample Output

clear ospf io-statistics on page 2077

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ospf io-statistics

```
user@host> clear ospf io-statistics
```

clear (ospf | ospf3) neighbor

List of Syntax

Syntax on page 2078
Syntax (EX Series Switch and QFX Series) on page 2078

Syntax

clear (ospf | ospf3) neighbor
<all>
<area *area-id*>
<instance *instance-name*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>
<neighbor>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>

Syntax (EX Series Switch and QFX Series)

clear (ospf | ospf3) neighbor
<all>
<area *area-id*>
<instance *instance-name*>
<interface *interface-name*>
<neighbor>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
realm option introduced in Junos OS Release 9.2.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Tear down Open Shortest Path First (OSPF) neighbor connections.

Options

all—Tear down OSPF connections with all neighbors for all routing instances.

area *area-id*—(Optional) Tear down neighbor connections for the specified area only.

instance *instance-name*—(Optional) Tear down neighbor connections for the specified routing instance only.

interface *interface-name*—(Optional) Tear down neighbor connections for the specified interface only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

neighbor—(Optional) Clear the state of the specified neighbor only.

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)—(Optional) (OSPFv3 only) Clear the state of the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

Required Privilege Level

clear

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- show (ospf | ospf3) neighbor on page 2133

List of Sample Output clear ospf neighbor all on page 2079

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ospf neighbor all

```
user@host> clear ospf neighbor all
```

clear (ospf | ospf3) overload

List of Syntax

Syntax on page 2080
Syntax (EX Series Switches) on page 2080

Syntax

clear (ospf | ospf3) overload
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

clear (ospf | ospf3) overload
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Clear the Open Shortest Path First (OSPF) overload bit and rebuild link-state advertisements (LSAs).

Options

none—Clear the overload bit and rebuild LSAs for all routing instances.

instance *instance-name*—(Optional) Clear the overload bit and rebuild LSAs for the specified routing instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

List of Sample Output

clear ospf overload on page 2080

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ospf overload

```
user@host> clear ospf overload
```

 Copyright © 2018, Juniper Networks, Inc.

clear (ospf | ospf3) statistics

List of Syntax
Syntax on page 2081
Syntax (EX Series Switch and QFX Series) on page 2081

Syntax
clear (ospf | osfp3) statistics
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>

Syntax (EX Series Switch and QFX Series)
clear (ospf | osfp3) statistics
<instance *instance-name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
realm option introduced in Junos OS Release 9.2.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description
Clear Open Shortest Path First (OSPF) statistics.

Options
none—Clear OSPF statistics.

instance *instance-name*—(Optional) Clear statistics for the specified routing instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)—(Optional) (OSPFv3 only) Clear statistics for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

Required Privilege Level
clear

Related Documentation
- show (ospf | ospf3) statistics on page 2151

List of Sample Output
clear ospf statistics on page 2082

Output Fields
See show (ospf | ospf3) statistics for an explanation of output fields.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

clear ospf statistics

The following sample output displays OSPF statistics before and after the **clear ospf statistics** command is entered:

user@host> **show ospf statistics**

```
Packet type             Total                  Last 5 seconds
                   Sent       Received        Sent        Received
   Hello           3254           2268           3               1
     DbD             41             46           0               0
   LSReq              8              7           0               0
LSUpdate            212            154           0               0
   LSAck             65             98           0               0

DBDs retransmitted     :                 3, last 5 seconds :          0
LSAs flooded           :                12, last 5 seconds :          0
LSAs flooded high-prio :                 0, last 5 seconds :          0
LSAs retransmitted     :                 0, last 5 seconds :          0
LSAs transmitted to nbr:                 3, last 5 seconds :          0
LSAs requested         :                 5, last 5 seconds :          0
LSAs acknowledged      :                19, last 5 seconds :          0

Flood queue depth      :               0
Total rexmit entries   :               0
db summaries           :               0
lsreq entries          :               0

Receive errors:
  626 subnet mismatches
```

user@host> **clear ospf statistics**

user@host> **show ospf statistics**

```
Packet type             Total                  Last 5 seconds
                   Sent       Received        Sent        Received
   Hello              3              1           3               1
     DbD              0              0           0               0
   LSReq              0              0           0               0
LSUpdate              0              0           0               0
   LSAck              0              0           0               0

DBDs retransmitted     :                 0, last 5 seconds :          0
LSAs flooded           :                 0, last 5 seconds :          0
LSAs flooded high-prio :                 0, last 5 seconds :          0
LSAs retransmitted     :                 0, last 5 seconds :          0
LSAs transmitted to nbr:                 0, last 5 seconds :          0
LSAs requested         :                 0, last 5 seconds :          0
LSAs acknowledged      :                 0, last 5 seconds :          0

Flood queue depth      :               0
Total rexmit entries   :               0
db summaries           :               0
lsreq entries          :               0

Receive errors:
  None
```


Copyright © 2018, Juniper Networks, Inc.

show (ospf | ospf3) backup coverage

Syntax

show (ospf | ospf3) backup coverage
<instance *instance-name*>
< logical-system (all | *logical-system-name*)>
<realm (ipv4-unicast | ipv46-unicast>
<topology *topology-name*>

Syntax (QFX Series)

show (ospf | ospf3) backup coverage
<instance *instance-name*>
<topology *topology-name*>

Release Information

Command introduced in Junos OS Release 10.0.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display information about the level of backup coverage available for all the nodes and prefixes in the network.

Options

none—Display information about the level backup coverage for all OSPF routing instances in all logical systems.

logical-system (all | *logical-system-name*)—(Optional) Display information about the level of backup coverage for all logical systems or for a specific logical system.

instance *instance-name*—(Optional) Display information about the level of backup coverage for a specific OSPF routing instance.

realm (ipv4-unicast | ipv6-unicast)—(Optional) (OSPFv3 only) Display information about the level of backup coverage for the specific OSPFv3 realm, or address family.

topology (default | *topology-name*)—(Optional) (OSPFv2 only) Display information about the level of backup coverage for the specific OSPF topology.

Required Privilege Level

view

Related Documentation

- show (ospf | ospf3) backup lsp on page 2087

List of Sample Output

show ospf backup coverage on page 2085
show ospf3 backup coverage on page 2085

Output Fields

Table 176 on page 2085 lists the output fields for the **show (ospf | ospf3) backup coverage** command. Output fields are listed in the approximate order in which they appear.

 Copyright © 2018, Juniper Networks, Inc.

Table 176: show (ospf | ospf3) backup coverage Output Fields

Field Name	Field Description
Node Coverage	Information about backup coverage for each OSPF node.
Area	Area number. Area **0.0.0.0** is the backbone.
Covered Nodes	Number of nodes for which backup coverage is available.
Total Nodes	Total number of OSPF nodes.
Route Coverage	Information about backup coverage for each type of OSPF route.
Path Type	Type of OSPF path: **Intra**, **Inter**, **Ext1**, **Ext2**, and **All**.
Covered Routes	For each path type, the number of routes for which backup coverage is available.
Total Routes	For each path type, the total number of configured routes.
Percent Covered	For all nodes and for each path type, the percentage for which backup coverage is available.

Sample Output

show ospf backup coverage

```
user@host> show ospf backup coverage
Topology default coverage:

Node Coverage:

Area                Covered  Total  Percent
                      Nodes  Nodes  Covered
0.0.0.0                   4      5   80.00%

Route Coverage:

Path Type  Covered   Total  Percent
            Routes  Routes  Covered
Intra            8      14   57.14%
Inter            0       0  100.00%
Ext1             0       0  100.00%
Ext2             1       1  100.00%
All              9      15   60.00%
```

show ospf3 backup coverage

```
user @host > show ospf3 backup coverage
show ospf3 backup coverage
Node Coverage:

Area                Covered  Total  Percent
                      Nodes  Nodes  Covered
```

```
0.0.0.0                4      5   80.00%

Route Coverage:

Path Type  Covered   Total  Percent
            Routes  Routes  Covered
Intra            4       6   66.67%
Inter            0       0  100.00%
Ext1             0       0  100.00%
Ext2             1       1  100.00%
All              5       7   71.43%
```

 Copyright © 2018, Juniper Networks, Inc.

show (ospf | ospf3) backup lsp

Syntax

show (ospf | ospf3) backup lsp
<logical-system (all | *logical-system-name*>
<realm (ipv4-unicast | ipv6-unicast)>

Release Information

Command introduced in Junos OS Release 10.0.

Description

Display information about MPLS label-switched-paths (LSPs) designated as backup routes for OSPF routes.

NOTE: MPLS LSPs can be used as backup routes only for routes in the default OSPFv2 topology and not for any configured topology. Additionally, MPLS LSPs cannot be used as backup routes for nondefault instances either for OSPFv2 or OSPFv3.

Options

none—Display information all MPLS LSPs designated as backup routes.

logical-system (all | ***logical-system-name*****)**—(Optional) Display information about MPLS LSPs designated as backup routes for all logical systems or a specific logical system.

realm (ipv4-unicast | ipv6-unicast)—(Optional) (OSPFv3 only) Display information about MPLS LSPs designated as backup routes for a specific realm, or address family.

Required Privilege Level

view

Related Documentation

- show (ospf | ospf3) backup coverage on page 2084

List of Sample Output

show ospf backup lsp on page 2088
show ospf3 backup lsp on page 2088

Output Fields

Table 177 on page 2087 lists the output fields for the **show (ospf | ospf3) backup lsp** command. Output fields are listed in the approximate order in which they appear.

Table 177: show (ospf | ospf3) backup lsp Output Fields

Field Name	Field Description
MPLS LSP name	Name of each MPLS LSP designated as a backup path.
Egress	IP address of the egress router for the LSP.

Copyright © 2018, Juniper Networks, Inc.

Table 177: show (ospf | ospf3) backup lsp Output Fields (continued)

Field Name	Field Description

Status	State of the LSP:
	• **Up**—The router can detect RSVP hello messages from the neighbor.
	• **Down**—The router has received one of the following indications:
	  ◦ Communication failure from the neighbor.
	  ◦ Communication from IGP that the neighbor is unavailable.
	  ◦ Change in the sequence numbers in the RSVP hello messages sent by the neighbor.
	• **Deleted**—The LSP is no longer available as a backup path.

Last change	Time elapsed since the neighbor state changed either from **up** or **down** or from **down** to **up**. The format is ***hh:mm:ss***.
TE-metric	Configured traffic engineering metric.
Metric	Configured metric.

Sample Output

show ospf backup lsp

```
user@host> show ospf backup lsp
tobanff
  Egress: 10.255.71.239, Status: up, Last change: 00:00:23
  TE-metric: 0, Metric: 0
```

Sample Output

show ospf3 backup lsp

```
user@host> show ospf3 backup lsp
tobanff
  Egress: 10.255.71.239, Status: up, Last change: 00:00:45
  TE-metric: 0, Metric: 0
```


Copyright © 2018, Juniper Networks, Inc.

show (ospf | ospf3) backup neighbor

Syntax

show (ospf | ospf3) backup neighbor
<area *area-id*>
<instance (default | *instance-name*)>
<logical-system (default | ipv4-multicast | *logical-system-name*)>
<topology (default | ipv4-multicast | *topology-name*)>

Syntax (QFX Series)

show (ospf | ospf3) backup neighbor
<area *area-id*>
<instance *instance-name*>
<topology (default | ipv4-multicast | *topology-name*)>

Release Information

Command introduced in Junos OS Release 10.0.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display the neighbors through which direct next hops for the backup paths are available.

Options

none—Display all neighbors that have direct next hops for backup paths.

area *area-id*—(Optional) Display the area information.

instance (default | *instance-name*)—(Optional) Display information about the default routing instance or a particular routing instance.

logical-system (default | ipv4-multicast | *logical-system-name*)—(Optional) Display information about the default logical system, IPv4 multicast logical system, or a particular logical system.

topology (default | ipv4-multicast | *topology-name*)—(OSPFv2 only) (Optional) Display information about the default topology, IPv4 multicast topology, or a particular topology.

Required Privilege Level

view

Related Documentation

- show (ospf | ospf3) backup spf on page 2091

List of Sample Output

show ospf backup neighbor on page 2090

Output Fields

Table 178 on page 2090 lists the output fields for the **show (ospf |ospf3) backup neighbor** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2018, Juniper Networks, Inc.

Table 178: show (ospf |ospf3) backup neighbor Output Fields

Field Name	Field Description	Level of Output
Neighbor to Self Metric	Metric from the backup neighbor to the OSPF node.	All levels
Self to Neighbor Metric	Metric from the OSPF node to the backup neighbor.	All levels
Direct next-hop	Interface and address of the direct next hop.	All levels

Sample Output

show ospf backup neighbor

```
user@host> show ospf backup neighbor
Topology default backup neighbors:

Area 0.0.0.5 backup neighbors:

10.0.0.5
  Neighbor to Self Metric: 5
  Self to Neighbor Metric: 5
  Direct next-hop: ge-4/0/0.111 via 10.0.175.5

10.0.0.6
  Neighbor to Self Metric: 5
  Self to Neighbor Metric: 5
  Direct next-hop: ge-4/1/0.110 via 10.0.176.6
```

 Copyright © 2018, Juniper Networks, Inc.

show (ospf | ospf3) backup spf

Syntax

show (ospf | ospf3) backup spf
<brief | detail>
<area *area-id*>
<instance *instance-name*>
<logical-system (all | *logical-system-name*>
<no-coverage>
<*node-id*>
<realm (ipv4–unicast | ipv6–unicast)>
<topology (default | ipv4-multicast | *topology-name*)>

Syntax (QFX Series)

show (ospf | ospf3) backup spf
<brief | detail>
<area *area-id*>
<instance *instance-name*>
<no-coverage>
<*node-id*>
<topology (default | ipv4-multicast | *topology-name*)>

Release Information

Command introduced in Junos OS Release 10.0.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

Display information about OSPF shortest-path-first calculations for backup paths.

Options

none—Display information about OSPF shortest-path-first (SPF) calculations for all backup paths for all destination nodes.

brief | detail—(Optional) Display the specified level of output.

area *area-id*—(Optional) Display the area information.

instance *instance-name*—(Optional) Display information about the routing instance.

logical-system (all | *logical-system-name*)—(Optional) Display information about all logical systems or a specific logical system.

no-coverage—(Optional) Display information if there is no backup coverage.

node-id—(Optional) Display information about the node specified.

realm (ipv4–unicast | ipv6–unicast)—(Optional) Display information about the **ipv4** or **ipv6** realm.

topology (default | ipv4-multicast | *topology-name*)—(Optional) (OSPFv2 only) Display information about the default topology, IPv4 multicast topology, or a specifc topology.

Required Privilege Level

view

Copyright © 2018, Juniper Networks, Inc.

List of Sample Output

Output Fields

Table 179 on page 2092 lists the output fields for the **show (ospf |ospf3) backup spf** command. Output fields are listed in the approximate order in which they appear.

Table 179: show (ospf |ospf3) backup spf Output Fields

Field Name	Field Description	Level of Output
Area *area-id* results	Area for which the results are displayed. Area 0.0.0.0 is the backbone area.	All levels
address	Address of the node for which the results are displayed.	All levels
Self to Destination Metric	Metric from the node to the destination.	All levels
Parent Node	Address of the parent node.	All levels
Primary next-hop	Address of the next hop.	All levels
Backup Neighbor	Address of the backup neighbor or LSP endpoint and the following information: • Neighbor to Destination Metric • Neighbor to Self Metric • Self to Neighbor Metric • Status (Eligible, Not Eligible, Not Evaluated) and the reason for the status. NOTE: If the backup neighbor is an LSP endpoint, it is indicated as such after the neighbor address.	All levels

Sample Output

show ospf backup spf

```
user@host> show ospf backup spf
Topology default results:

Area 0.0.0.0 results:

pro16-d-lo0.xxx.yyyy.net
  Self to Destination Metric: 1
  Parent Node: pro16-b-lo0.xxx.yyyy.net
  Primary next-hop: at-1/0/1.0
  Backup Neighbor: pro16-c-lo0.xxx.yyyy.net (LSP endpoint)
    Neighbor to Destination Metric: 4, Neighbor to Self Metric: 3
    Self to Neighbor Metric: 3
    Not eligible, Reason: Path loops
  Backup Neighbor: pro16-d-lo0.xxx.yyyy.net
    Neighbor to Destination Metric: 0, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing
  ...
```


Copyright © 2018, Juniper Networks, Inc.

show ospf backup spf detail

```
user@host> show ospf backup spf detail
Topology default results:

Area 0.0.0.0 results:

11.14.10.2
  Self to Destination Metric: 1
  Parent Node: 10.255.70.103
  Primary next-hop: ae0.0
  Backup Neighbor: 10.255.71.243
    Neighbor to Destination Metric: 2, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops
  Backup Neighbor: 10.255.71.52
    Neighbor to Destination Metric: 15, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing
  Backup Neighbor: 10.255.71.242
    Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops

10.255.71.52
  Self to Destination Metric: 1
  Parent Node: 11.14.10.2
  Primary next-hop: ae0.0 via 11.14.10.2
  Backup Neighbor: 10.255.71.52
    Neighbor to Destination Metric: 0, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing
  Backup Neighbor: 10.255.71.243
    Neighbor to Destination Metric: 2, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops
  Backup Neighbor: 10.255.71.242
    Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops

10.255.71.242
  Self to Destination Metric: 1
  Parent Node: 10.255.70.103
  Primary next-hop: as0.0
  Backup Neighbor: 10.255.71.242
    Neighbor to Destination Metric: 0, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing
  Backup Neighbor: 10.255.71.243
    Neighbor to Destination Metric: 2, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops
  Backup Neighbor: 10.255.71.52
    Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops

10.255.71.243
  Self to Destination Metric: 1
```

Copyright © 2018, Juniper Networks, Inc.

```
    Parent Node: 10.255.70.103
    Primary next-hop: so-6/0/0.0
    Backup Neighbor: 10.255.71.243
      Neighbor to Destination Metric: 0, Neighbor to Self Metric: 1
      Self to Neighbor Metric: 1
      Not eligible, Reason: Primary next-hop link fate sharing
    Backup Neighbor: 10.255.71.52
      Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
      Self to Neighbor Metric: 1
      Not eligible, Reason: Path loops
    Backup Neighbor: 10.255.71.242
      Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
      Self to Neighbor Metric: 1
      Not eligible, Reason: Path loops

  12.15.0.1
    Self to Destination Metric: 2
    Parent Node: 10.255.71.243
    Primary next-hop: so-6/0/0.0
    Backup next-hop: ae0.0 via 11.14.10.2
    Backup Neighbor: 10.255.71.243
      Neighbor to Destination Metric: 1, Neighbor to Self Metric: 1
      Self to Neighbor Metric: 1
      Not eligible, Reason: Primary next-hop link fate sharing
    Backup Neighbor: 10.255.71.52
      Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
      Self to Neighbor Metric: 1
      Eligible, Reason: Contributes backup next-hop
    Backup Neighbor: 10.255.71.242
      Neighbor to Destination Metric: 17, Neighbor to Self Metric: 15
      Self to Neighbor Metric: 1
      Not evaluated, Reason: Interface is already covered

  10.255.71.238
    Self to Destination Metric: 2
    Parent Node: 10.255.71.243
    Primary next-hop: so-6/0/0.0
    Backup next-hop: as0.0
    Backup Neighbor: 10.255.71.243
      Neighbor to Destination Metric: 1, Neighbor to Self Metric: 1
      Self to Neighbor Metric: 1
      Not eligible, Reason: Primary next-hop link fate sharing
    Backup Neighbor: 10.255.71.242
      Neighbor to Destination Metric: 15, Neighbor to Self Metric: 15
      Self to Neighbor Metric: 1
      Eligible, Reason: Contributes backup next-hop
    Backup Neighbor: 10.255.71.52
      Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
      Self to Neighbor Metric: 1
      Not evaluated, Reason: Interface is already covered

  10.255.71.239
    Self to Destination Metric: 2
    Parent Node: 12.15.0.1
    Primary next-hop: so-6/0/0.0
    Backup next-hop: ae0.0 via 11.14.10.2
    Backup Neighbor: 10.255.71.243
      Neighbor to Destination Metric: 1, Neighbor to Self Metric: 1
      Self to Neighbor Metric: 1
      Not eligible, Reason: Primary next-hop link fate sharing
    Backup Neighbor: 10.255.71.52
```

 Copyright © 2018, Juniper Networks, Inc.

```
    Neighbor to Destination Metric: 15, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Eligible, Reason: Contributes backup next-hop
  Backup Neighbor: 10.255.71.242
    Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not evaluated, Reason: Interface is already covered

14.15.0.2
  Self to Destination Metric: 3
  Parent Node: 10.255.71.239
  Primary next-hop: so-6/0/0.0
  Backup next-hop: ae0.0 via 11.14.10.2
  Backup Neighbor: 10.255.71.243
    Neighbor to Destination Metric: 2, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing
  Backup Neighbor: 10.255.71.52
    Neighbor to Destination Metric: 15, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Eligible, Reason: Contributes backup next-hop
  Backup Neighbor: 10.255.71.242
    Neighbor to Destination Metric: 17, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not evaluated, Reason: Interface is already covered
```

show ospf3 backup spf

```
user@host> show ospf3 backup spf
Area 0.0.0.0 results:

10.255.71.52;0.0.0.5
  Self to Destination Metric: 1
  Parent Node: 10.255.70.103
  Primary next-hop: ae0.0
  Backup Neighbor: 10.255.71.243
    Neighbor to Destination Metric: 2, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops
  Backup Neighbor: 10.255.71.52
    Neighbor to Destination Metric: 15, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing
  Backup Neighbor: 10.255.71.242
    Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops

10.255.71.52
  Self to Destination Metric: 1
  Parent Node: 10.255.71.52;0.0.0.5
  Primary next-hop: ae0.0 via fe80::290:69ff:fe0f:67f0
  Backup Neighbor: 10.255.71.52
    Neighbor to Destination Metric: 0, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing
  Backup Neighbor: 10.255.71.243
    Neighbor to Destination Metric: 2, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops
  Backup Neighbor: 10.255.71.242
```

Copyright © 2018, Juniper Networks, Inc.

```
    Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops

10.255.71.242
  Self to Destination Metric: 1
  Parent Node: 10.255.70.103
  Primary next-hop: as0.0
  Backup Neighbor: 10.255.71.242
    Neighbor to Destination Metric: 0, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing
  Backup Neighbor: 10.255.71.243
    Neighbor to Destination Metric: 2, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops
  Backup Neighbor: 10.255.71.52
    Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops

10.255.71.243
  Self to Destination Metric: 1
  Parent Node: 10.255.70.103
  Primary next-hop: so-6/0/0.0
  Backup Neighbor: 10.255.71.243
    Neighbor to Destination Metric: 0, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing
  Backup Neighbor: 10.255.71.52
    Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops
  Backup Neighbor: 10.255.71.242
    Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops

10.255.71.243;0.0.0.2
  Self to Destination Metric: 2
  Parent Node: 10.255.71.243
  Primary next-hop: so-6/0/0.0
  Backup next-hop: ae0.0 via fe80::290:69ff:fe0f:67f0
  Backup Neighbor: 10.255.71.243
    Neighbor to Destination Metric: 1, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing
  Backup Neighbor: 10.255.71.52
    Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Eligible, Reason: Contributes backup next-hop
  Backup Neighbor: 10.255.71.242
    Neighbor to Destination Metric: 17, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not evaluated, Reason: Interface is already covered

10.255.71.238
  Self to Destination Metric: 2
  Parent Node: 10.255.71.243
  Primary next-hop: so-6/0/0.0
  Backup next-hop: as0.0
```

 Copyright © 2018, Juniper Networks, Inc.

```
  Backup Neighbor: 10.255.71.243
    Neighbor to Destination Metric: 1, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing
  Backup Neighbor: 10.255.71.242
    Neighbor to Destination Metric: 15, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Eligible, Reason: Contributes backup next-hop
  Backup Neighbor: 10.255.71.52
    Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not evaluated, Reason: Interface is already covered

10.255.71.239
  Self to Destination Metric: 2
  Parent Node: 10.255.71.243;0.0.0.2
  Primary next-hop: so-6/0/0.0
  Backup next-hop: ae0.0 via fe80::290:69ff:fe0f:67f0
  Backup Neighbor: 10.255.71.243
    Neighbor to Destination Metric: 1, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing
  Backup Neighbor: 10.255.71.52
    Neighbor to Destination Metric: 15, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Eligible, Reason: Contributes backup next-hop
  Backup Neighbor: 10.255.71.242
    Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not evaluated, Reason: Interface is already covered

10.255.71.239;0.0.0.4
  Self to Destination Metric: 3
  Parent Node: 10.255.71.239
  Primary next-hop: so-6/0/0.0
  Backup next-hop: ae0.0 via fe80::290:69ff:fe0f:67f0
  Backup Neighbor: 10.255.71.243
    Neighbor to Destination Metric: 2, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing
  Backup Neighbor: 10.255.71.52
    Neighbor to Destination Metric: 15, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Eligible, Reason: Contributes backup next-hop
  Backup Neighbor: 10.255.71.242
    Neighbor to Destination Metric: 17, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not evaluated, Reason: Interface is already covered
```

Copyright © 2018, Juniper Networks, Inc.

show ospf context-identifier

List of Syntax

Syntax on page 2098
Syntax (EX Series Switches and QFX Series) on page 2098

Syntax

show ospf context-identifier
<brief | detail>
<area *area-id*>
<context-id>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches and QFX Series)

show ospf context-identifier
<brief | detail>
<area *area-id*>
<context-id>
<instance *instance-name*>

Release Information

Command introduced in Junos OS Release 10.4.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display the context identifier information processed and advertised by Open Shortest Path First (OSPF) for egress protection.

Options

none—Display information about all context identifiers.

brief | detail—(Optional) Display the specified level of output.

area *area-id*—(Optional) Display information about the context identifier for the specified area.

context-id—(Optional) Display information about the specified context identifier.

instance *instance-name*—(Optional) Display information about the context identifier for the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- *egress-protection (Layer 2 circuit)* in the *Junos OS VPNs Library for Routing Devices*
- *egress-protection (MPLS)* in the *Junos OS VPNs Library for Routing Devices*

List of Sample Output

show ospf context-identifier on page 2099
show ospf context-identifier detail on page 2099

Copyright © 2018, Juniper Networks, Inc.

Output Fields Table 180 on page 2099 lists the output fields for the **show ospf context-identifier** command. Output fields are listed in the approximate order in which they appear.

Table 180: show ospf context-identifier Output Fields

Field Name	Field Description	Level of Output
Context	IPv4 address that defines a protection pair. The context is manually configured on both primary and protector provider edge (PE) devices.	All levels
Status	State of the path: **active** or **inactive**.	All levels
Metric	Advertised OSPF metric.	All levels
Area	OSPF area number.	All levels
Other Advertisements	Other advertisements received by the OSPF node: • **Advertising router**– Address of the device that sent the advertisement. • **Type**–Type of OSPF path: **inter-area** and **stub**. • **Metric**–Advertised OSPF metric. • **None**–No additional advertisements were received by the OSPF node.	**detail**

Sample Output

show ospf context-identifier

```
user@host> show ospf context-identifier
Context-id: 2.2.4.3
  Status: active, Metric: 65534, PE role: protector, Area: 0.0.0.0
```

show ospf context-identifier detail

```
user@host> show ospf context-identifier detail
Context-id: 88.24.13.1
                Status: inactive, Metric: 0, PE role: protector, Area: 0.0.0.13
                Other Advertisements:
                  Advertising router: 8.8.8.103
                    Type: stub link
                    Metric: 65534
```

show ospf database

List of Syntax

Syntax on page 2100
Syntax (EX Series Switches and QFX Series) on page 2100

Syntax

show ospf database
<brief | detail | extensive | summary>
<advertising-router (*address* | self)>
<area *area-id*>
<asbrsummary>
<external>
<instance *instance-name*>
<link-local>
<logical-system (all | *logical-system-name*)>
<lsa-id *lsa-id*>
<netsummary>
<network>
<nssa>
<opaque-area>
<router>

Syntax (EX Series Switches and QFX Series)

show ospf database
<brief | detail | extensive | summary>
<advertising-router (*address* | self)>
<area *area-id*>
<asbrsummary>
<external>
<instance *instance-name*>
<link-local>
<lsa-id *lsa-id*>
<netsummary>
<network>
<nssa>
<opaque-area>
<router>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
advertising-router self (***address*** | **self**) option introduced in Junos OS Release 9.5.
advertising-router self (***address*** | **self**) option introduced in Junos OS Release 9.5 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display the entries in the OSPF version 2 (OSPFv2) link-state database, which contains data about link-state advertisement (LSA) packets.

Options

none—Display standard information about entries in the OSPFv2 link-state database for all routing instances.

brief | detail | extensive | summary—(Optional) Display the specified level of output.

 Copyright © 2018, Juniper Networks, Inc.

advertising-router (*address* | self)—(Optional) Display the LSAs advertised either by a particular routing device or by this routing device.

area *area-id*—(Optional) Display the LSAs in a particular area.

asbrsummary—(Optional) Display summary AS boundary router LSA entries.

external—(Optional) Display external LSAs.

instance *instance-name*—(Optional) Display all OSPF database information under the named routing instance.

link-local—(Optional) Display information about link-local LSAs.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsa-id *lsa-id*—(Optional) Display the LSA with the specified LSA identifier.

netsummary—(Optional) Display summary network LSAs.

network—(Optional) Display information about network LSAs.

nssa—(Optional) Display information about not-so-stubby area (NSSA) LSAs.

opaque-area—(Optional) Display opaque area-scope LSAs.

router—(Optional) Display information about router LSAs.

Required Privilege Level view

Related Documentation

- clear (ospf | ospf3) database on page 2074

List of Sample Output

show ospf database on page 2103
show ospf database on page 2104
show ospf database brief on page 2104
show ospf database detail on page 2104
show ospf database extensive on page 2105
show ospf database summary on page 2108
show ospf database opaque-area detail on page 2108

Output Fields Table 181 on page 2101 describes the output fields for the **show ospf database** command. Output fields are listed in the approximate order in which they appear.

Table 181: show ospf database Output Fields

Field Name	Field Description	Level of Output
area	Area number. Area **0.0.0.0** is the backbone area.	All levels

Table 181: show ospf database Output Fields (continued)

Field Name	Field Description	Level of Output
Type	Type of link advertisement: **ASBRSum**, **Extern**, **Network**, **NSSA**, **OpaqArea**, **Router**, or **Summary**.	All levels
ID	LSA identifier included in the advertisement. An asterisk preceding the identifier marks database entries that originated from the local routing device.	All levels
Adv Rtr	Address of the routing device that sent the advertisement.	All levels
Seq	Link sequence number of the advertisement.	All levels
Age	Time elapsed since the LSA was originated, in seconds.	All levels
Opt	Optional OSPF capabilities associated with the LSA.	All levels
Cksum	Checksum value of the LSA.	All levels
Len	Length of the advertisement, in bytes.	All levels
Router	Router link-state advertisement information: • **bits**—Flags describing the routing device that generated the LSP. • **link count**—Number of links in the advertisement. • **id**—ID of a routing device or subnet on the link. • **data**—For stub networks, the subnet mask. Otherwise, the IP address of the routing device that generated the LSP. • **type**—Type of link. It can be **PointToPoint**, **Transit**, **Stub**, or **Virtual**. • **TOS count**—Number of type-of-service (ToS) entries in the advertisement. • **TOS 0 metric**—Metric for ToS 0. • **TOS**—Type-of-service (ToS) value. • **metric**—Metric for the ToS.	**detail extensive**
Network	Network link-state advertisement information: • **mask**—Network mask. • **attached router**—ID of the attached neighbor.	**detail extensive**
Summary	Summary link-state advertisement information: • **mask**—Network mask. • **TOS**—Type-of-service (ToS) value. • **metric**—Metric for the ToS.	**detail extensive**
Gen timer	How long until the LSA is regenerated.	**extensive**
Aging timer	How long until the LSA expires.	**extensive**
Installed ***hh:mm:ss*** **ago**	How long ago the route was installed.	**extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 181: show ospf database Output Fields (continued)

Field Name	Field Description	Level of Output
expires in *hh:mm:ss*	How long until the route expires.	**extensive**
sent *hh:mm:ss* ago	How long ago the LSA was sent.	**extensive**
Last changed *hh:mm:ss* ago	How long ago the route was changed.	**extensive**
Change count	Number of times the route has changed.	**extensive**
Ours	Indicates that this is a local advertisement.	**extensive**
Router LSAs	Number of router link-state advertisements in the link-state database.	**summary**
Network LSAs	Number of network link-state advertisements in the link-state database.	**summary**
Summary LSAs	Number of summary link-state advertisements in the link-state database.	**summary**
NSSA LSAs	Number of not-so-stubby area link-state advertisements in the link-state database.	**summary**

Sample Output

show ospf database

```
user@host> show ospf database
OSPF link state database, Area 0.0.0.1
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   10.255.70.103    10.255.70.103    0x80000002   215  0x20 0x4112  48
Router  *10.255.71.242    10.255.71.242    0x80000002   214  0x20 0x11b1  48
Summary *23.1.1.0         10.255.71.242    0x80000002   172  0x20 0x6d72  28
Summary *24.1.1.0         10.255.71.242    0x80000002   177  0x20 0x607e  28
NSSA    *33.1.1.1         10.255.71.242    0x80000002   217  0x28 0x73bd  36

    OSPF link state database, Area 0.0.0.2
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   10.255.71.52     10.255.71.52     0x80000004   174  0x20 0xd021  36
Router  *10.255.71.242    10.255.71.242    0x80000003   173  0x20 0xe191  36
Network *23.1.1.1         10.255.71.242    0x80000002   173  0x20 0x9c76  32
Summary *12.1.1.0         10.255.71.242    0x80000001   217  0x20 0xfeec  28
Summary *24.1.1.0         10.255.71.242    0x80000002   177  0x20 0x607e  28
NSSA    *33.1.1.1         10.255.71.242    0x80000001   222  0x28 0xe047  36

    OSPF link state database, Area 0.0.0.3
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   10.255.71.238    10.255.71.238    0x80000003   179  0x20 0x3942  36
Router  *10.255.71.242    10.255.71.242    0x80000003   177  0x20 0xf37d  36
Network *24.1.1.1         10.255.71.242    0x80000002   177  0x20 0xc591  32
Summary *12.1.1.0         10.255.71.242    0x80000001   217  0x20 0xfeec  28
Summary *23.1.1.0         10.255.71.242    0x80000002   172  0x20 0x6d72  28
NSSA    *33.1.1.1         10.255.71.242    0x80000001   222  0x28 0xeb3b  36
```

Copyright © 2018, Juniper Networks, Inc.

show ospf database

The output for **show ospf databse nssa** with **nssa-only** configuration statement enabled at **[edit policy-options policy-statement *policy-name* term *term name* then external]**, which clears P-bit on type 7 LSA.

```
user@host> show ospf database
OSPF link state database, Area 0.0.0.1
 Type       ID               Adv Rtr          Seq      Age  Opt  Cksum  Len
Router   10.255.70.103    10.255.70.103    0x80000002  215  0x20 0x4112  48
Router  *10.255.71.242    10.255.71.242    0x80000002  214  0x20 0x11b1  48
Summary *23.1.1.0         10.255.71.242    0x80000002  172  0x20 0x6d72  28
Summary *24.1.1.0         10.255.71.242    0x80000002  177  0x20 0x607e  28
NSSA    *33.1.1.1         10.255.71.242    0x80000002  217  0x20 0x73bd  36

    OSPF link state database, Area 0.0.0.2
 Type       ID               Adv Rtr          Seq      Age  Opt  Cksum  Len
Router   10.255.71.52     10.255.71.52     0x80000004  174  0x20 0xd021  36
Router  *10.255.71.242    10.255.71.242    0x80000003  173  0x20 0xe191  36
Network *23.1.1.1         10.255.71.242    0x80000002  173  0x20 0x9c76  32
Summary *12.1.1.0         10.255.71.242    0x80000001  217  0x20 0xfeec  28
Summary *24.1.1.0         10.255.71.242    0x80000002  177  0x20 0x607e  28
NSSA    *33.1.1.1         10.255.71.242    0x80000001  222  0x28 0xe047  36

    OSPF link state database, Area 0.0.0.3
 Type       ID               Adv Rtr          Seq      Age  Opt  Cksum  Len
Router   10.255.71.238    10.255.71.238    0x80000003  179  0x20 0x3942  36
Router  *10.255.71.242    10.255.71.242    0x80000003  177  0x20 0xf37d  36
Network *24.1.1.1         10.255.71.242    0x80000002  177  0x20 0xc591  32
Summary *12.1.1.0         10.255.71.242    0x80000001  217  0x20 0xfeec  28
Summary *23.1.1.0         10.255.71.242    0x80000002  172  0x20 0x6d72  28
NSSA    *33.1.1.1         10.255.71.242    0x80000001  222  0x20 0xeb3b  36
```

show ospf database brief

The output for the **show ospf database brief** command is identical to that for the **show ospf database** command. For sample output, see show ospf database on page 2103.

show ospf database detail

```
user@host> show ospf database detail
    OSPF link state database, Area 0.0.0.1
 Type       ID               Adv Rtr          Seq      Age  Opt  Cksum  Len
Router   10.255.70.103    10.255.70.103    0x80000002  261  0x20 0x4112  48
  bits 0x0, link count 2
  id 10.255.71.242, data 12.1.1.1, Type PointToPoint (1)
  TOS count 0, TOS 0 metric 1
  id 12.1.1.0, data 255.255.255.0, Type Stub (3)
  TOS count 0, TOS 0 metric 1
Router  *10.255.71.242    10.255.71.242    0x80000002  260  0x20 0x11b1  48
  bits 0x3, link count 2
  id 10.255.70.103, data 12.1.1.2, Type PointToPoint (1)
  TOS count 0, TOS 0 metric 1
  id 12.1.1.0, data 255.255.255.0, Type Stub (3)
  TOS count 0, TOS 0 metric 1
Summary *23.1.1.0         10.255.71.242    0x80000002  218  0x20 0x6d72  28
  mask 255.255.255.0
  TOS 0x0, metric 1
```

 Copyright © 2018, Juniper Networks, Inc.

```
Summary *24.1.1.0        10.255.71.242    0x80000002   223  0x20 0x607e  28
  mask 255.255.255.0
  TOS 0x0, metric 1
NSSA    *33.1.1.1        10.255.71.242    0x80000002   263  0x28 0x73bd  36
  mask 255.255.255.255
  Type 2, TOS 0x0, metric 0, fwd addr 12.1.1.2, tag 0.0.0.0

    OSPF link state database, Area 0.0.0.2
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   10.255.71.52     10.255.71.52     0x80000004   220  0x20 0xd021  36
  bits 0x0, link count 1
  id 23.1.1.1, data 23.1.1.2, Type Transit (2)
  TOS count 0, TOS 0 metric 1
Router  *10.255.71.242    10.255.71.242    0x80000003   219  0x20 0xe191  36
  bits 0x3, link count 1
  id 23.1.1.1, data 23.1.1.1, Type Transit (2)
  TOS count 0, TOS 0 metric 1
Network *23.1.1.1        10.255.71.242    0x80000002   219  0x20 0x9c76  32
  mask 255.255.255.0
  attached router 10.255.71.242
  attached router 10.255.71.52
Summary *12.1.1.0        10.255.71.242    0x80000001   263  0x20 0xfeec  28
  mask 255.255.255.0
  TOS 0x0, metric 1
Summary *24.1.1.0        10.255.71.242    0x80000002   223  0x20 0x607e  28
  mask 255.255.255.0
  TOS 0x0, metric 1
NSSA    *33.1.1.1        10.255.71.242    0x80000001   268  0x28 0xe047  36
  mask 255.255.255.255
  Type 2, TOS 0x0, metric 0, fwd addr 23.1.1.1, tag 0.0.0.0

    OSPF link state database, Area 0.0.0.3
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   10.255.71.238    10.255.71.238    0x80000003   225  0x20 0x3942  36
  bits 0x0, link count 1
  id 24.1.1.1, data 24.1.1.2, Type Transit (2)
  TOS count 0, TOS 0 metric 1
Router  *10.255.71.242    10.255.71.242    0x80000003   223  0x20 0xf37d  36
  bits 0x3, link count 1
  id 24.1.1.1, data 24.1.1.1, Type Transit (2)
  TOS count 0, TOS 0 metric 1
Network *24.1.1.1        10.255.71.242    0x80000002   223  0x20 0xc591  32
  mask 255.255.255.0
  attached router 10.255.71.242
  attached router 10.255.71.238
Summary *12.1.1.0        10.255.71.242    0x80000001   263  0x20 0xfeec  28
  mask 255.255.255.0
  TOS 0x0, metric 1
Summary *23.1.1.0        10.255.71.242    0x80000002   218  0x20 0x6d72  28
  mask 255.255.255.0
  TOS 0x0, metric 1
NSSA    *33.1.1.1        10.255.71.242    0x80000001   268  0x28 0xeb3b  36
  mask 255.255.255.255
  Type 2, TOS 0x0, metric 0, fwd addr 24.1.1.1, tag 0.0.0.0
```

show ospf database extensive

```
user@host> show ospf database extensive
    OSPF link state database, Area 0.0.0.1
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   10.255.70.103    10.255.70.103    0x80000002   286  0x20 0x4112  48
```

```
  bits 0x0, link count 2
  id 10.255.71.242, data 12.1.1.1, Type PointToPoint (1)
  TOS count 0, TOS 0 metric 1
  id 12.1.1.0, data 255.255.255.0, Type Stub (3)
  TOS count 0, TOS 0 metric 1
  Aging timer 00:55:14
  Installed 00:04:43 ago, expires in 00:55:14
  Last changed 00:04:43 ago, Change count: 2
Router  *10.255.71.242    10.255.71.242    0x80000002   285  0x20 0x11b1  48
  bits 0x3, link count 2
  id 10.255.70.103, data 12.1.1.2, Type PointToPoint (1)
  TOS count 0, TOS 0 metric 1
  id 12.1.1.0, data 255.255.255.0, Type Stub (3)
  TOS count 0, TOS 0 metric 1
  Gen timer 00:45:15
  Aging timer 00:55:15
  Installed 00:04:45 ago, expires in 00:55:15, sent 00:04:43 ago
  Last changed 00:04:45 ago, Change count: 2, Ours
Summary *23.1.1.0         10.255.71.242    0x80000002   243  0x20 0x6d72  28
  mask 255.255.255.0
  TOS 0x0, metric 1
  Gen timer 00:45:57
  Aging timer 00:55:57
  Installed 00:04:03 ago, expires in 00:55:57, sent 00:04:01 ago
  Last changed 00:04:48 ago, Change count: 1, Ours
Summary *24.1.1.0         10.255.71.242    0x80000002   248  0x20 0x607e  28
  mask 255.255.255.0
  TOS 0x0, metric 1
  Gen timer 00:45:52
  Aging timer 00:55:52
  Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
  Last changed 00:04:48 ago, Change count: 1, Ours
NSSA    *33.1.1.1         10.255.71.242    0x80000002   288  0x28 0x73bd  36
  mask 255.255.255.255
  Type 2, TOS 0x0, metric 0, fwd addr 12.1.1.2, tag 0.0.0.0
  Gen timer 00:45:12
  Aging timer 00:55:12
  Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:48 ago
  Last changed 00:04:48 ago, Change count: 2, Ours

    OSPF link state database, Area 0.0.0.2
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   10.255.71.52     10.255.71.52     0x80000004   245  0x20 0xd021  36
  bits 0x0, link count 1
  id 23.1.1.1, data 23.1.1.2, Type Transit (2)
  TOS count 0, TOS 0 metric 1
  Aging timer 00:55:55
  Installed 00:04:02 ago, expires in 00:55:55
  Last changed 00:04:02 ago, Change count: 2
Router  *10.255.71.242    10.255.71.242    0x80000003   244  0x20 0xe191  36
  bits 0x3, link count 1
  id 23.1.1.1, data 23.1.1.1, Type Transit (2)
  TOS count 0, TOS 0 metric 1
  Gen timer 00:45:56
  Aging timer 00:55:56
  Installed 00:04:04 ago, expires in 00:55:56, sent 00:04:02 ago
  Last changed 00:04:04 ago, Change count: 2, Ours
Network *23.1.1.1         10.255.71.242    0x80000002   244  0x20 0x9c76  32
  mask 255.255.255.0
  attached router 10.255.71.242
  attached router 10.255.71.52
```

 Copyright © 2018, Juniper Networks, Inc.

```
    Gen timer 00:45:56
    Aging timer 00:55:56
    Installed 00:04:04 ago, expires in 00:55:56, sent 00:04:02 ago
    Last changed 00:04:04 ago, Change count: 1, Ours
  Summary *12.1.1.0         10.255.71.242    0x80000001   288  0x20 0xfeec  28
    mask 255.255.255.0
    TOS 0x0, metric 1
    Gen timer 00:45:12
    Aging timer 00:55:12
    Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:04 ago
    Last changed 00:04:48 ago, Change count: 1, Ours
  Summary *24.1.1.0         10.255.71.242    0x80000002   248  0x20 0x607e  28
    mask 255.255.255.0
    TOS 0x0, metric 1
    Gen timer 00:45:52
    Aging timer 00:55:52
    Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:04 ago
    Last changed 00:04:48 ago, Change count: 1, Ours
  NSSA    *33.1.1.1         10.255.71.242    0x80000001   293  0x28 0xe047  36
    mask 255.255.255.255
    Type 2, TOS 0x0, metric 0, fwd addr 23.1.1.1, tag 0.0.0.0
    Gen timer 00:45:07
    Aging timer 00:55:07
    Installed 00:04:53 ago, expires in 00:55:07, sent 00:04:04 ago
    Last changed 00:04:53 ago, Change count: 1, Ours

      OSPF link state database, Area 0.0.0.3
   Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
  Router   10.255.71.238    10.255.71.238    0x80000003   250  0x20 0x3942  36
    bits 0x0, link count 1
    id 24.1.1.1, data 24.1.1.2, Type Transit (2)
    TOS count 0, TOS 0 metric 1
    Aging timer 00:55:50
    Installed 00:04:07 ago, expires in 00:55:50
    Last changed 00:04:07 ago, Change count: 2
  Router  *10.255.71.242    10.255.71.242    0x80000003   248  0x20 0xf37d  36
    bits 0x3, link count 1
    id 24.1.1.1, data 24.1.1.1, Type Transit (2)
    TOS count 0, TOS 0 metric 1
    Gen timer 00:45:52
    Aging timer 00:55:52
    Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
    Last changed 00:04:08 ago, Change count: 2, Ours
  Network *24.1.1.1         10.255.71.242    0x80000002   248  0x20 0xc591  32
    mask 255.255.255.0
    attached router 10.255.71.242
    attached router 10.255.71.238
    Gen timer 00:45:52
    Aging timer 00:55:52
    Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
    Last changed 00:04:08 ago, Change count: 1, Ours
  Summary *12.1.1.0         10.255.71.242    0x80000001   288  0x20 0xfeec  28
    mask 255.255.255.0
    TOS 0x0, metric 1
    Gen timer 00:45:12
    Aging timer 00:55:12
    Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:13 ago
    Last changed 00:04:48 ago, Change count: 1, Ours
  Summary *23.1.1.0         10.255.71.242    0x80000002   243  0x20 0x6d72  28
    mask 255.255.255.0
    TOS 0x0, metric 1
```

```
  Gen timer 00:45:57
  Aging timer 00:55:57
  Installed 00:04:03 ago, expires in 00:55:57, sent 00:04:01 ago
  Last changed 00:04:48 ago, Change count: 1, Ours
NSSA    *33.1.1.1        10.255.71.242    0x80000001   293  0x28 0xeb3b  36
  mask 255.255.255.255
  Type 2, TOS 0x0, metric 0, fwd addr 24.1.1.1, tag 0.0.0.0
  Gen timer 00:45:07
  Aging timer 00:55:07
  Installed 00:04:53 ago, expires in 00:55:07, sent 00:04:13 ago
  Last changed 00:04:53 ago, Change count: 1, Ours
```

show ospf database summary

```
user@host> show ospf database summary
Area 0.0.0.1:
   2 Router LSAs
   2 Summary LSAs
   1 NSSA LSAs
Area 0.0.0.2:
   2 Router LSAs
   1 Network LSAs
   2 Summary LSAs
   1 NSSA LSAs
Area 0.0.0.3:
   2 Router LSAs
   1 Network LSAs
   2 Summary LSAs
   1 NSSA LSAs
Externals:
Interface fe-2/2/1.0:
Interface ge-0/3/2.0:
Interface so-0/1/2.0:
Interface so-0/1/2.0:
```

show ospf database opaque-area detail

```
user@host> show ospf database opaque-area detail
    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
OpaqArea*4.0.0.0         1.1.1.10         0x800000ef  1510  0x22 0x2810 156
Area-opaque LSA
  SR-Algorithm (8), length 1:
    Algo (1), length 1:
        0
  SID/Label Range (9), length 12:
    Range Size (1), length 3:
        256
    SID/Label (1), length 3:
      Label (1), length 3:
        802048
  SID/Label Range (9), length 12:
    Range Size (1), length 3:
   SID/Label (1), length 3:
      Label (1), length 3:
        802304
  SID/Label Range (9), length 12:
    Range Size (1), length 3:
        256
    SID/Label (1), length 3:
```

 Copyright © 2018, Juniper Networks, Inc.

```
      Label (1), length 3:
        802560
  SID/Label Range (9), length 12:
    Range Size (1), length 3:
        256
    SID/Label (1), length 3:
      Label (1), length 3:
        802816
  SID/Label Range (9), length 12:
    Range Size (1), length 3:
        256
    SID/Label (1), length 3:

      Label (1), length 3:
        803072
  SID/Label Range (9), length 12:
    Range Size (1), length 3:
        256
    SID/Label (1), length 3:
      Label (1), length 3:
        803328
  SID/Label Range (9), length 12:
    Range Size (1), length 3:
        256
    SID/Label (1), length 3:
      Label (1), length 3:
        803584
  SID/Label Range (9), length 12:
    Range Size (1), length 3:
        256
    SID/Label (1), length 3:
      Label (1), length 3:
        803840

The Extended Prefix LSA (eg):

Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
OpaqArea*7.0.0.1         10.10.10.10      0x80000002   561  0x22 0x60eb  44
Area-opaque LSA
  Extended Prefix (1), length 20:
      Route Type (1), length 1:
        1
      Prefix Length (2), length 1:
        32
      AF (3), length 1:
        0
      Flags (4),  length 1:
        0x40
      Prefix (5), length 32:
        10.10.10.10
      Prefix Sid (2), length 8:
       Flags (1), length 1:
        0x00
       MT ID (2), length 1:
        0
       Algorithm (3), length 1:
        0
       SID (4), length 4:
        0

Extended Links LSA (eg):
```

```
Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
OpaqArea 8.0.0.1          1.1.1.1          0x80000001   688  0x22 0xcd8a  48
Area-opaque LSA
  Extended Link (1), length 24:
    Link Type (1), length 1:
     2
     Link Id (2), length 4:
     11.1.1.1
     Link Data (3), length 4:
     11.1.1.2
    Adjacency Sid (2), length 7:
        Flags (1), length 1:
          0x60
        MT ID (2), length 1:
           0
        Weight (3), length 1:
           0
        Label (4), length 3:
          300416
```

 Copyright © 2018, Juniper Networks, Inc.

show ospf3 database

List of Syntax

Syntax on page 2111
Syntax (EX Series Switches and QFX Series) on page 2111

Syntax

show ospf3 database
<brief | detail | extensive | summary>
<advertising-router (*address* | self)>
<area *area-id*>
<external>
<instance *instance-name*>
<inter-area-prefix>
<inter-area-router>
<intra-area-prefix>
<link>
<link-local>
<logical-system (all | *logical-system-name*)>
<lsa-id *lsa-id*>
<network>
<nssa>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>
<router>

Syntax (EX Series Switches and QFX Series)

show ospf3 database
<brief | detail | extensive | summary>
<advertising-router (*address* | self)>
<area *area-id*>
<external>
<instance *instance-name*>
<inter-area-prefix>
<inter-area-router>
<intra-area-prefix>
<link>
<link-local>
<lsa-id *lsa-id*>
<network>
<nssa>
<router>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
realm option introduced in Junos OS Release 9.2.
advertising-router (***address*** | **self**) option introduced in Junos Relase 9.5.
advertising-router (***address*** | **self**) option introduced in Junos OS Release 9.5 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description

Display the entries in the OSPF version 3 (OSPFv3) link-state database, which contains data about link-state advertisement (LSA) packets.

Options

none—Display standard information about all entries in the OSPFv3 link-state database.

Copyright © 2018, Juniper Networks, Inc.

brief | detail | extensive | summary—(Optional) Display the specified level of output.

advertising-router (*address* | self)—(Optional) Display the LSAs advertised either by a particular routing device or by this routing device.

area *area-id*—(Optional) Display the LSAs in a particular area.

external—(Optional) Display external LSAs.

instance *instance-name*—(Optional) Display all OSPF database information under the named routing instance.

inter-area-prefix—(Optional) Display information about interarea-prefix LSAs.

inter-area-router—(Optional) Display information about interarea-router LSAs.

intra-area-prefix—(Optional) Display information about intra-area-prefix LSAs.

link—(Optional) Display information about link LSAs.

link-local—(Optional) Display information about link-local LSAs.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsa-id *lsa-id*—(Optional) Display the LSA with the specified LSA identifier.

network—(Optional) Display information about network LSAs.

nssa—(Optional) Display information about not-so-stubby area (NSSA) LSAs.

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)—(Optional) Display information about the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family other than IPv6 unicast, which is the default.

router—(Optional) Display information about router LSAs.

Required Privilege Level

view

Related Documentation

- clear (ospf | ospf3) database on page 2074

List of Sample Output

show ospf3 database brief on page 2117
show ospf3 database extensive on page 2118
show ospf3 database summary on page 2120

Output Fields

Table 182 on page 2113 lists the output fields for the **show ospf3 database** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2018, Juniper Networks, Inc.

Table 182: show ospf3 database Output Fields

Field Name	Field Description	Level of Output
OSPF link state database, area ***area-number***	Entries in the link-state database for this area.	**brief detail extensive**
OSPF AS SCOPE link state database	Entries in the AS scope link-state database.	**brief detail extensive**
OSPF Link-Local link state database, interface ***interface-name***	Entries in the link-local link-state database for this interface.	**brief detail extensive**
area	Area number. Area **0.0.0.0** is the backbone area.	All levels
Type	Type of link advertisement: **Extern**, **InterArPfx**, **InterArRtr**, **IntraArPrx** , **Link**, **Network**, **NSSA**, or **Router**.	**brief detail extensive**
ID	Link identifier included in the advertisement. An asterisk (*) preceding the identifier marks database entries that originated from the local routing device.	**brief detail extensive**
Adv Rtr	Address of the routing device that sent the advertisement.	**brief detail extensive**
Seq	Link sequence number of the advertisement.	**brief detail extensive**
Age	Time elapsed since the LSA was originated, in seconds.	**brief detail extensive**
Cksum	Checksum value of the LSA.	**brief detail extensive**
Len	Length of the advertisement, in bytes.	**brief detail extensive**
Router (Router Link-State Advertisements)		
bits	Flags describing the routing device that generated the LSP.	**detail extensive**
Options	Option bits carried in the router LSA.	**detail extensive**
For Each Router Link		
Type	Type of interface. The value of all other output fields describing a routing device interface depends on the interface's type: • **PointToPoint (1)**—Point-to-point connection to another routing device. • **Transit (2)**—Connection to a transit network. • **Virtual (4)**—Virtual link.	**detail extensive**
Loc-if-id	Local interface ID assigned to the interface that uniquely identifies the interface with the routing device.	**detail extensive**
Nbr-if-id	Interface ID of the neighbor's interface for this routing device link.	**detail extensive**

Table 182: show ospf3 database Output Fields (continued)

Field Name	Field Description	Level of Output
Nbr-rtr-id	Router ID of the neighbor routing device (for type 2 interfaces, the attached link's designated router).	**detail extensive**
Metric	Cost of the router link.	**detail extensive**
Gen timer	How long until the LSA is regenerated, in the format ***hours*:*minutes*:*seconds***.	**extensive**
Aging timer	How long until the LSA expires, in the format ***hours*:*minutes*:*seconds***.	**extensive**
Installed *nn*:*nn*:*nn* ago	How long ago the route was installed, in the format ***hours*:*minutes*:*seconds***.	**extensive**
expires in *nn*:*nn*:*nn*	How long until the route expires, in the format ***hours*:*minutes*:*seconds***.	**extensive**
sent *nn*:*nn*:*nn* ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format ***hours*:*minutes*:*seconds***.	**extensive**
Ours	Indicates that this is a local advertisement.	**extensive**
Network (Network Link-State Advertisements)		
Options	Option bits carried in the network LSA.	**detail extensive**
Attached Router	Router IDs of each of the routing devices attached to the link. Only routing devices that are fully adjacent to the designated router are listed. The designated router includes itself in this list.	**detail extensive**
InterArPfx (Interarea-Prefix Link-State Advertisements)		
Prefix	IPv6 address prefix.	**detail extensive**
Prefix-options	Option bit associated with the prefix.	**detail extensive**
Metric	Cost of this route. Expressed in the same units as the interface costs in the router LSAs. When the interarea-prefix LSA is describing a route to a range of addresses, the cost is set to the maximum cost to any reachable component of the address range.	**detail extensive**
Gen timer	How long until the LSA is regenerated, in the format ***hours*:*minutes*:*seconds***.	**extensive**
Aging timer	How long until the LSA expires, in the format ***hours*:*minutes*:*seconds***.	**extensive**
Installed *nn*:*nn*:*nn* ago	How long ago the route was installed, in the format ***hours*:*minutes*:*seconds***.	**extensive**
expires in *nn*:*nn*:*nn*	How long until the route expires, in the format ***hours*:*minutes*:*seconds***.	**extensive**
sent *nn*:*nn*:*nn* ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format ***hours*:*minutes*:*seconds***.	**extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 182: show ospf3 database Output Fields (continued)

Field Name	Field Description	Level of Output
Ours	Indicates that this is a local advertisement.	**extensive**
InterArRtr (Interarea-Router Link-State Advertisements)		
Dest-router-id	Router ID of the routing device described by the LSA.	**detail extensive**
options	Optional capabilities supported by the routing device.	**detail extensive**
Metric	Cost of this route. Expressed in the same units as the interface costs in the router LSAs. When the interarea-prefix LSA is describing a route to a range of addresses, the cost is set to the maximum cost to any reachable component of the address range.	**detail extensive**
Prefix	IPv6 address prefix.	**extensive**
Prefix-options	Option bit associated with the prefix.	**extensive**
Extern (External Link-State Advertisements)		
Prefix	IPv6 address prefix.	**detail extensive**
Prefix-options	Option bit associated with the prefix.	**detail extensive**
Metric	Cost of the route, which depends on the value of **Type**.	**detail extensive**
Type ***n***	Type of external metric: **Type 1** or **Type 2**.	**detail extensive**
Aging timer	How long until the LSA expires, in the format ***hours*:*minutes*:*seconds***.	**extensive**
Installed ***nn:nn:nn*** **ago**	How long ago the route was installed, in the format ***hours*:*minutes*:*seconds***.	**extensive**
expires in ***nn:nn:nn***	How long until the route expires, in the format ***hours*:*minutes*:*seconds***.	**extensive**
sent ***nn:nn:nn*** **ago**	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format ***hours*:*minutes*:*seconds***.	**extensive**
Link (Link-State Advertisements)		
IPv6-Address	IPv6 link-local address on the link for which this link LSA originated.	**detail extensive**
Options	Option bits carried in the link LSA.	**detail extensive**
priority	Router priority of the interface attaching the originating routing device to the link.	**detail extensive**
Prefix-count	Number of IPv6 address prefixes contained in the LSA. The rest of the link LSA contains a list of IPv6 prefixes to be associated with the link.	**detail extensive**
Prefix	IPv6 address prefix.	**detail extensive**

Table 182: show ospf3 database Output Fields (continued)

Field Name	Field Description	Level of Output
Prefix-options	Option bit associated with the prefix.	**detail extensive**
Gen timer	How long until the LSA is regenerated, in the format ***hours*:*minutes*:*seconds***.	**extensive**
Aging timer	How long until the LSA expires, in the format ***hours*:*minutes*:*seconds***.	**extensive**
Installed *nn:nn:nn* ago	How long ago the route was installed, in the format ***hours*:*minutes*:*seconds***.	**extensive**
expires in *nn:nn:nn*	How long until the route expires, in the format ***hours*:*minutes*:*seconds***.	**extensive**
sent *nn:nn:nn* ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format ***hours*:*minutes*:*seconds***.	**extensive**
Ours	Indicates that this is a local advertisement.	**extensive**
IntraArPfx (Intra-Area-Prefix Link-State Advertisements)		
Ref-lsa-type	LSA type of the referenced LSA. • **Router**—Address prefixes are associated with a router LSA. • **Network**—Address prefixes are associated with a network LSA.	**detail extensive**
Ref-lsa-id	Link-state ID of the referenced LSA.	**detail extensive**
Ref-router-id	Advertising router ID of the referenced LSA.	**detail extensive**
Prefix-count	Number of IPv6 address prefixes contained in the LSA. The rest of the link LSA contains a list of IPv6 prefixes to be associated with the link.	**detail extensive**
Prefix	IPv6 address prefix.	**detail extensive**
Prefix-options	Option bit associated with the prefix.	**detail extensive**
Metric	Cost of this prefix. Expressed in the same units as the interface costs in the router LSAs.	**detail extensive**
Gen timer	How long until the LSA is regenerated, in the format ***hours*:*minutes*:*seconds***.	**extensive**
Aging timer	How long until the LSA expires, in the format ***hours*:*minutes*:*seconds***.	**extensive**
Installed *hh:mm:ss* ago	How long ago the route was installed, in the format ***hours*:*minutes*:*seconds***.	**extensive**
expires in *hh:mm:ss*	How long until the route expires, in the format ***hours*:*minutes*:*seconds***.	**extensive**
sent *hh:mm:ss* ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format ***hours*:*minutes*:*seconds***.	**extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 182: show ospf3 database Output Fields (continued)

Field Name	Field Description	Level of Output
n Router LSAs	Number of router LSAs in the link-state database.	summary
n Network LSAs	Number of network LSAs in the link-state database.	summary
n InterArPfx LSAs	Number of interarea-prefix LSAs in the link-state database.	summary
n InterArRtr LSAs	Number of interarea-router LSAs in the link-state database.	summary
n IntraArPfx LSAs	Number of intra-area-prefix LSAs in the link-state database.	summary
Externals	Display of the external LSA database.	summary
n Extern LSAs	Number of external LSAs in the link-state database.	summary
Interface *interface-name*	Name of the interface for which link-local LSA information is displayed.	summary
n Link LSAs	Number of link LSAs in the link-state database.	summary

Sample Output

show ospf3 database brief

```
user@host> show ospf3 database brief
    OSPF3 link state database, area 0.0.0.0
 Type        ID               Adv Rtr          Seq         Age  Cksum  Len
Router       0.0.0.1          10.255.4.85      0x80000003  885  0xa697  40
Router      *0.0.0.1          10.255.4.93      0x80000002  953  0xc677  40
InterArPfx *0.0.0.2          10.255.4.93      0x80000001  910  0xb96f  44
InterArRtr *0.0.0.1          10.255.4.93      0x80000001  910  0xe159  32
IntraArPfx *0.0.0.1          10.255.4.93      0x80000002  432  0x788f  72

    OSPF3 link state database, area 0.0.0.1
 Type        ID               Adv Rtr          Seq         Age  Cksum  Len
Router      *0.0.0.1          10.255.4.93      0x80000003  916  0xea40  40
Router       0.0.0.1          10.255.4.97      0x80000006  851  0xc95b  40
Network      0.0.0.2          10.255.4.97      0x80000002  916  0x4598  32
InterArPfx *0.0.0.1          10.255.4.93      0x80000002  117  0xa980  44
InterArPfx *0.0.0.2          10.255.4.93      0x80000002   62  0xd47e  44
NSSA         0.0.0.1          10.255.4.97      0x80000002  362  0x45ee  44
IntraArPfx  0.0.0.1          10.255.4.97      0x80000006  851  0x2f77  52

    OSPF3 AS SCOPE link state database
 Type        ID               Adv Rtr          Seq         Age  Cksum  Len
Extern       0.0.0.1          10.255.4.85      0x80000002   63  0x9b86  44
Extern      *0.0.0.1          10.255.4.93      0x80000001  910  0x59c9  44

    OSPF3 Link-Local link state database, interface ge-1/3/0.0
 Type        ID               Adv Rtr          Seq         Age  Cksum  Len
Link        *0.0.0.2          10.255.4.93      0x80000003  916  0x4dab  64
```

show ospf3 database extensive

```
user@host> show ospf3 database extensive
    OSPF3 link state database, area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Router      0.0.0.1          10.255.4.85      0x80000003  1028  0xa697  40
  bits 0x2, Options 0x13
  Type PointToPoint (1), Metric 10
      Loc-If-Id 2, Nbr-If-Id 3, Nbr-Rtr-Id 10.255.4.93
  Aging timer 00:42:51
  Installed 00:17:05 ago, expires in 00:42:52, sent 02:37:54 ago
Router     *0.0.0.1          10.255.4.93      0x80000002  1096  0xc677  40
  bits 0x3, Options 0x13
  Type PointToPoint (1), Metric 10
      Loc-If-Id 3, Nbr-If-Id 2, Nbr-Rtr-Id 10.255.4.85
  Gen timer 00:00:40
  Aging timer 00:41:44
  Installed 00:18:16 ago, expires in 00:41:44, sent 00:18:14 ago
  Ours
InterArPfx *0.0.0.2          10.255.4.93      0x80000001  1053  0xb96f  44
  Prefix feee::10:10:2:0/126
  Prefix-options 0x0, Metric 10
  Gen timer 00:17:02
  Aging timer 00:42:26
  Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago
  Ours
InterArPfx *0.0.0.3          10.255.4.93      0x80000001  1053  0x71d3  44
  Prefix feee::10:255:4:97/128
  Prefix-options 0x0, Metric 10
  Gen timer 00:21:07
  Aging timer 00:42:26
  Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago
  Ours
InterArRtr *0.0.0.1          10.255.4.93      0x80000001  1053  0xe159  32
  Dest-router-id 10.255.4.97, Options 0x19, Metric 10
  Gen timer 00:29:18
  Aging timer 00:42:26
  Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago
  Ours
IntraArPfx  0.0.0.1          10.255.4.85      0x80000002  1028  0x2403  72
  Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.255.4.85
  Prefix-count 2
  Prefix feee::10:255:4:85/128
    Prefix-options 0x2, Metric 0
  Prefix feee::10:10:1:0/126
    Prefix-options 0x0, Metric 10
  Aging timer 00:42:51
  Installed 00:17:05 ago, expires in 00:42:52, sent 02:37:54 ago
IntraArPfx *0.0.0.1          10.255.4.93      0x80000002   575  0x788f  72
  Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.255.4.93
  Prefix-count 2
  Prefix feee::10:255:4:93/128
    Prefix-options 0x2, Metric 0
  Prefix feee::10:10:1:0/126
    Prefix-options 0x0, Metric 10
  Gen timer 00:33:23
  Aging timer 00:50:24
  Installed 00:09:35 ago, expires in 00:50:25, sent 00:09:33 ago
    OSPF3 link state database, area 0.0.0.1
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Router     *0.0.0.1          10.255.4.93      0x80000003  1059  0xea40  40
```


Copyright © 2018, Juniper Networks, Inc.

```
    bits 0x3, Options 0x19
    Type Transit (2), Metric 10
        Loc-If-Id 2, Nbr-If-Id 2, Nbr-Rtr-Id 10.255.4.97
    Gen timer 00:08:51
    Aging timer 00:42:20
    Installed 00:17:39 ago, expires in 00:42:21, sent 00:17:37 ago
  Router      0.0.0.1          10.255.4.97      0x80000006   994  0xc95b  40
    bits 0x2, Options 0x19
    Type Transit (2), Metric 10
        Loc-If-Id 2, Nbr-If-Id 2, Nbr-Rtr-Id 10.255.4.97
    Aging timer 00:43:25
    Installed 00:16:31 ago, expires in 00:43:26, sent 02:37:54 ago
  Network     0.0.0.2          10.255.4.97      0x80000002  1059  0x4598  32
    Options 0x11
    Attached router 10.255.4.97
    Attached router 10.255.4.93
    Aging timer 00:42:20
    Installed 00:17:36 ago, expires in 00:42:21, sent 02:37:54 ago
  InterArPfx *0.0.0.1          10.255.4.93      0x80000002   260  0xa980  44
    Prefix feee::10:10:1:0/126
    Prefix-options 0x0, Metric 10
    Gen timer 00:45:39
    Aging timer 00:55:39
    Installed 00:04:20 ago, expires in 00:55:40, sent 00:04:18 ago
    Ours
  InterArPfx *0.0.0.2          10.255.4.93      0x80000002   205  0xd47e  44
    Prefix feee::10:255:4:93/128
    Prefix-options 0x0, Metric 0
    Gen timer 00:46:35
    Aging timer 00:56:35
    Installed 00:03:25 ago, expires in 00:56:35, sent 00:03:23 ago
    Ours
  InterArPfx *0.0.0.3          10.255.4.93      0x80000001  1089  0x9bbb  44
    Prefix feee::10:255:4:85/128
    Prefix-options 0x0, Metric 10
    Gen timer 00:04:46
    Aging timer 00:41:51
    Installed 00:18:09 ago, expires in 00:41:51, sent 00:17:43 ago
    Ours
  NSSA        0.0.0.1          10.255.4.97      0x80000002   505  0x45ee  44
    Prefix feee::200:200:1:0/124
    Prefix-options 0x8, Metric 10, Type 2,
    Aging timer 00:51:35
    Installed 00:08:22 ago, expires in 00:51:35, sent 02:37:54 ago
  IntraArPfx  0.0.0.1          10.255.4.97      0x80000006   994  0x2f77  52
    Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.255.4.97
    Prefix-count 1
    Prefix feee::10:255:4:97/128
      Prefix-options 0x2, Metric 0
    Aging timer 00:43:25
    Installed 00:16:31 ago, expires in 00:43:26, sent 02:37:54 ago
  IntraArPfx  0.0.0.3          10.255.4.97      0x80000002  1059  0x4446  52
    Ref-lsa-type Network, Ref-lsa-id 0.0.0.2, Ref-router-id 10.255.4.97
    Prefix-count 1
    Prefix feee::10:10:2:0/126
      Prefix-options 0x0, Metric 0
    Aging timer 00:42:20
    Installed 00:17:36 ago, expires in 00:42:21, sent 02:37:54 ago
      OSPF3 AS SCOPE link state database
   Type       ID               Adv Rtr           Seq         Age  Cksum  Len
  Extern      0.0.0.1          10.255.4.85      0x80000002   206  0x9b86  44
```

```
  Prefix feee::100:100:1:0/124
  Prefix-options 0x0, Metric 20, Type 2,
  Aging timer 00:56:34
  Installed 00:03:23 ago, expires in 00:56:34, sent 02:37:54 ago
Extern     *0.0.0.1          10.255.4.93      0x80000001  1053  0x59c9  44
  Prefix feee::200:200:1:0/124
  Prefix-options 0x0, Metric 10, Type 2,
  Gen timer 00:25:12
  Aging timer 00:42:26
  Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago

    OSPF3 Link-Local link state database, interface ge-1/3/0.0
 Type       ID                Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.2          10.255.4.93      0x80000003  1059  0x4dab  64
  fe80::290:69ff:fe39:1cdb
  Options 0x11, priority 128
  Prefix-count 1
  Prefix feee::10:10:2:0/126 Prefix-options 0x0
  Gen timer 00:12:56
  Aging timer 00:42:20
  Installed 00:17:39 ago, expires in 00:42:21, sent 00:17:37 ago
Link        0.0.0.2          10.255.4.97      0x80000003   205  0xa87d  64
  fe80::290:69ff:fe38:883e
  Options 0x11, priority 128
  Prefix-count 1
  Prefix feee::10:10:2:0/126 Prefix-options 0x0
  Aging timer 00:56:35
  Installed 00:03:22 ago, expires in 00:56:35, sent 02:37:54 ago

    OSPF3 Link-Local link state database, interface so-2/2/0.0
 Type       ID                Adv Rtr           Seq         Age  Cksum  Len
Link        0.0.0.2          10.255.4.85      0x80000002   506  0x42bb  64
  fe80::280:42ff:fe10:f169
  Options 0x13, priority 128
  Prefix-count 1
  Prefix feee::10:10:1:0/126 Prefix-options 0x0
  Aging timer 00:51:34
  Installed 00:08:23 ago, expires in 00:51:34, sent 02:37:54 ago
Link       *0.0.0.3          10.255.4.93      0x80000002   505  0x6b7a  64
  fe80::280:42ff:fe10:f177
  Options 0x13, priority 128
  Prefix-count 1
  Prefix feee::10:10:1:0/126 Prefix-options 0x0
  Gen timer 00:37:28
  Aging timer 00:51:35
  Installed 00:08:25 ago, expires in 00:51:35, sent 00:08:23 ago
  Ours
```

show ospf3 database summary

```
user@host> show ospf3 database summary
Area 0.0.0.0:
   2 Router LSAs
   1 InterArPfx LSAs
   1 InterArRtr LSAs
   1 IntraArPfx LSAs
Area 0.0.0.1:
   2 Router LSAs
   1 Network LSAs
   2 InterArPfx LSAs
   1 NSSA LSAs
```

 Copyright © 2018, Juniper Networks, Inc.

```
   1 IntraArPfx LSAs
Externals:
   2 Extern LSAs
Interface ge-1/3/0.0:
   1 Link LSAs
Interface lo0.0:
Interface so-2/2/0.0:
   1 Link LSAs
```

show (ospf | ospf3) interface

List of Syntax

Syntax on page 2122
Syntax (EX Series Switches and QFX Series) on page 2122

Syntax

show (ospf | ospf3) interface
<brief | detail | extensive>
<area *area-id*>
<*interface-name*>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<realm (ip4-multicast | ipv4-unicast | ipv6-multicast)>

Syntax (EX Series Switches and QFX Series)

show (ospf | ospf3) interface
<brief | detail | extensive>
<area *area-id*>
<*interface-name*>
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
area option introduced in Junos OS Release 9.2.
area option introduced in Junos OS Release 9.2 for EX Series switches.
realm option introduced in Junos OS Release 9.2.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display the status of OSPF interfaces.

Options

none—Display standard information about the status of all OSPF interfaces for all routing instances

brief | detail | extensive—(Optional) Display the specified level of output.

area *area-id*—(Optional) Display information about the interfaces that belong to the specified area.

interface-name—(Optional) Display information for the specified interface.

instance *instance-name*—(Optional) Display all OSPF interfaces under the named routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)—(OSPFv3 only) (Optional) Display information about the interfaces for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

 Copyright © 2018, Juniper Networks, Inc.

Required Privilege Level view

List of Sample Output

show ospf interface brief on page 2125
show ospf interface detail on page 2125
show ospf3 interface detail on page 2125
show ospf interface detail(When Multiarea Adjacency Is Configured) on page 2125
show ospf interface area area-id on page 2127
show ospf interface extensive (When Flooding Reduction Is Enabled) on page 2127
show ospf interface extensive (When LDP Synchronization Is Configured) on page 2127

Output Fields Table 183 on page 2123 lists the output fields for the **show (ospf | ospf3) interface** command. Output fields are listed in the approximate order in which they appear.

Table 183: show (ospf | ospf3) interface Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface running OSPF version 2 or OSPF version 3.	All levels
State	State of the interface: **BDR**, **Down**, **DR**, **DRother**, **Loop**, **PtToPt**, or **Waiting**.	All levels
Area	Number of the area that the interface is in.	All levels
DR ID	Address of the area's designated router.	All levels
BDR ID	Backup designated router for a particular subnet.	All levels
Nbrs	Number of neighbors on this interface.	All levels
Type	Type of interface: **LAN**, **NBMA**, **P2MP**, **P2P**, or **Virtual**.	**detail extensive**
Address	IP address of the neighbor.	**detail extensive**
Mask	Netmask of the neighbor.	**detail extensive**
Prefix-length	(OSPFv3) IPv6 prefix length, in bits.	**detail extensive**
OSPF3-Intf-Index	(OSPFv3) OSPF version 3 interface index.	**detail extensive**
MTU	Interface maximum transmission unit (MTU).	**detail extensive**
Cost	Interface cost (metric).	**detail extensive**
DR addr	Address of the designated router.	**detail extensive**
BDR addr	Address of the backup designated router.	**detail extensive**
Adj count	Number of adjacent neighbors.	**detail extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 183: show (ospf | ospf3) interface Output Fields (continued)

Field Name	Field Description	Level of Output
Secondary	Indicates that this interface is configured as a secondary interface for this area. This interface can belong to more than one area, but can be designated as a primary interface for only one area.	**detail extensive**
Flood Reduction	Indicates that this interface is configured with flooding reduction. All self-originated LSAs from this interface are initially sent with the **DoNotAge** bit set. As a result, LSAs are refreshed only when a change occurs.	**extensive**
Priority	Router priority used in designated router (DR) election on this interface.	**detail extensive**
Flood list	List of link-state advertisements (LSAs) that might be about to flood this interface.	**extensive**
Ack list	Acknowledgment list. List of pending acknowledgments on this interface.	**extensive**
Descriptor list	List of packet descriptors.	**extensive**
Hello	Configured value for the hello timer.	**detail extensive**
Dead	Configured value for the dead timer.	**detail extensive**
Auth type	(OSPFv2) Authentication mechanism for sending and receiving OSPF protocol packets: • **MD5**—The MD5 mechanism is configured in accordance with RFC 2328. • **None**—No authentication method is configured. • **Password**—A simple password (RFC 2328) is configured.	**detail extensive**
Topology	(Multiarea adjacency) Name of topology: **default** or ***name***.	
LDP sync state	(OSPFv2 and LDP synchronization) Current state of LDP synchronization: **in sync**, **in holddown**, and **not supported**.	**extensive**
reason	(OSPFv2 and LDP synchronization) Reason for the current state of LDP synchronization. The LDP session might be up or down, or adjacency might be up or down.	**extensive**
config holdtime	(OSPFv2 and LDP synchronization) Configured value of the hold timer. If the state is not synchronized, and the hold time is not infinity, the **remaining** field displays the number of seconds that remain until the configured hold timer expires.	**extensive**
IPSec SA name	(OSPFv2) Name of the IPSec security association name.	**detail extensive**
Active key ID	(OSPFv2 and MD5) Number from **0** to **255** that uniquely identifies an MD5 key.	**detail extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 183: show (ospf | ospf3) interface Output Fields (continued)

Field Name	Field Description	Level of Output
Start time	(OSPFv2 and MD5) Time at which the routing device starts using an MD5 key to authenticate OSPF packets transmitted on the interface on which this key is configured. To authenticate received OSPF protocol packets, the key becomes effective immediately after the configuration is committed. If the start time option is not configured, the key is effective immediately for send and receive and is displayed as **Start time 1970 Jan 01 00:00:00 PST.**	**detail extensive**
ReXmit	Configured value for the Retransmit timer.	**detail extensive**
Stub, **Not Stub**, or **Stub NSSA**	Type of area.	**detail extensive**

Sample Output

show ospf interface brief

```
user@host>  show ospf interface brief
Intf                State    Area            DR ID           BDR ID         Nbrs
at-5/1/0.0          PtToPt  0.0.0.0         0.0.0.0         0.0.0.0           1
ge-2/3/0.0          DR      0.0.0.0         192.168.4.16    192.168.4.15      1
lo0.0               DR      0.0.0.0         192.168.4.16    0.0.0.0           0
so-0/0/0.0          Down    0.0.0.0         0.0.0.0         0.0.0.0           0
so-6/0/1.0          PtToPt  0.0.0.0         0.0.0.0         0.0.0.0           1
so-6/0/2.0          Down    0.0.0.0         0.0.0.0         0.0.0.0           0
so-6/0/3.0          PtToPt  0.0.0.0         0.0.0.0         0.0.0.0           1
```

show ospf interface detail

```
user@host>  show ospf interface detail
Interface               State     Area            DR ID           BDR ID Nbrs
fe-0/0/1.0              BDR      0.0.0.0         192.168.37.12   10.255.245.215 1
Type LAN, address 192.168.37.11, Mask 255.255.255.248, MTU 4460, Cost 40
DR addr 192.168.37.12, BDR addr 192.168.37.11, Adj count 1, Priority 128
Hello 10, Dead 40, ReXmit 5, Not Stub
t1-0/2/1.0              PtToPt   0.0.0.0         0.0.0.0         0.0.0.0 0
Type P2P, Address 0.0.0.0, Mask 0.0.0.0, MTU 1500, Cost 2604
  Adj count 0
Hello 10, Dead 40, ReXmit 5, Not Stub
Auth type: MD5, Active key ID 3, Start time 2002 Nov 19 10:00:00 PST
IPsec SA Name: sa
```

show ospf3 interface detail

```
user@host> show ospf3 interface so-0/0/3.0 detail
Interface               State     Area           DR-ID          BDR-ID      Nbrs
so-0/0/3.0              PtToPt   0.0.0.0        0.0.0.0        0.0.0.0        1
Address fe80::2a0:a5ff:fe28:1dfc, Prefix-length 64
OSPF3-Intf-index 1, Type P2P, MTU 4470, Cost 12, Adj-count 1
Hello 10, Dead 40, ReXmit 5, Not Stub
```

show ospf interface detail

Copyright © 2018, Juniper Networks, Inc.

(When Multiarea Adjacency Is Configured)

user@host> **show ospf interface detail**

```
user@host> show ospf interface detail
Interface           State   Area            DR ID           BDR ID          Nbrs
lo0.0               DR      0.0.0.0         10.255.245.2    0.0.0.0            0

  Type: LAN, Address: 127.0.0.1, Mask: 255.255.255.255, MTU: 65535, Cost: 0
  DR addr: 127.0.0.1, Adj count: 0, Priority: 128
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Cost: 0
lo0.0               DR      0.0.0.0         10.255.245.2    0.0.0.0            0

  Type: LAN, Address: 10.255.245.2, Mask: 255.255.255.255, MTU: 65535, Cost: 0
  DR addr: 10.255.245.2, Adj count: 0, Priority: 128
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Cost: 0
so-0/0/0.0          PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1

  Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Cost: 1
so-0/0/0.0          PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            0

  Type: P2P, Address: 192.168.37.46, Mask: 255.255.255.254, MTU: 4470, Cost: 1
  Adj count: 0, , Passive
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Passive, Cost: 1
so-1/0/0.0          PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1

  Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Cost: 1
so-1/0/0.0          PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            0

  Type: P2P, Address: 192.168.37.54, Mask: 255.255.255.254, MTU: 4470, Cost: 1
  Adj count: 0, , Passive
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Passive, Cost: 1
so-0/0/0.0          PtToPt  1.1.1.1         0.0.0.0         0.0.0.0            1

  Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
  Adj count: 1, Secondary
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Cost: 1
so-1/0/0.0          PtToPt  1.1.1.1         0.0.0.0         0.0.0.0            1

  Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
  Adj count: 1, Secondary
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Cost: 1
so-0/0/0.0          PtToPt  2.2.2.2         0.0.0.0         0.0.0.0            1
```

 Copyright © 2018, Juniper Networks, Inc.

```
  Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
  Adj count: 1, Secondary
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Cost: 1
so-1/0/0.0          PtToPt  2.2.2.2         0.0.0.0         0.0.0.0            1

  Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
  Adj count: 1, Secondary
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Cost: 1
```

show ospf interface area area-id

```
user@host> show ospf interface area 1.1.1.1
Interface           State   Area            DR ID           BDR ID          Nbrs
so-0/0/0.0          PtToPt  1.1.1.1         0.0.0.0         0.0.0.0            1
so-1/0/0.0          PtToPt  1.1.1.1         0.0.0.0         0.0.0.0            1
```

show ospf interface extensive (When Flooding Reduction Is Enabled)

```
user@host> show ospf interface extensive
Interface           State   Area            DR ID           BDR ID          Nbrs
fe-0/0/0.0          PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            0

  Type: P2P, Address: 10.10.10.1, Mask: 255.255.255.0, MTU: 1500, Cost: 1
  Adj count: 0
  Secondary, Flood Reduction
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Cost: 1
```

show ospf interface extensive (When LDP Synchronization Is Configured)

```
user@host> show ospf interface extensive
Interface              State     Area             DR ID            BDR ID
Nbrs
so-1/0/3.0             Down     0.0.0.0          0.0.0.0          0.0.0.0
0
  Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 65535
  Adj count: 0
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  LDP sync state: in holddown, for: 00:00:08, reason: LDP down during config
           config holdtime: 10 seconds, remaining: 1
```

show (ospf | ospf3) io-statistics

List of Syntax

Syntax on page 2128
Syntax (EX Series Switch and QFX Series) on page 2128

Syntax

```
show (ospf | ospf3) io-statistics
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and QFX Series)

```
show (ospf | ospf3) io-statistics
```

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display Open Shortest Path First (OSPF) input and output statistics.

Options

none—Display OSPF input and output statistics.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- clear (ospf | ospf3) statistics on page 2081

List of Sample Output

show ospf io-statistics on page 2129

Output Fields

Table 184 on page 2128 lists the output fields for the **show ospf io-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 184: show (ospf | ospf3) io-statistics Output Fields

Field Name	Field Description
Packets read	Number of OSPF packets read since the last time the routing protocol was started.
average per run	Total number of packets divided by the total number of times the OSPF read operation is scheduled to run.
max run	Maximum number of packets for a given run among all scheduled runs.
Receive errors	Number of faulty packets received with errors.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show ospf io-statistics

user@host> **show ospf io-statistics**

```
Packets read: 7361, average per run: 1.00, max run: 1
Receive errors:
  None
```

show (ospf | ospf3) log

List of Syntax

Syntax on page 2130
Syntax (EX Series Switch and QFX Series) on page 2130

Syntax

show (ospf | osfp3) log
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>
<topology *topology-name*>

Syntax (EX Series Switch and QFX Series)

show (ospf | osfp3) log
<instance *instance-name*>
<topology *topology-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
topology option introduced in Junos OS Release 9.0.
topology option introduced in Junos OS Release 9.0 for EX Series switches.
realm option introduced in Junos OS Release 9.2.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display the entries in the Open Shortest Path First (OSPF) log of SPF calculations.

Options

none—Display entries in the OSPF log of SPF calculations for all routing instances.

instance *instance-name*—(Optional) Display entries for the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

topology *topology-name*—(Optional) (OSPFv2 only) Display entries for the specified topology.

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)—(OSPFv3 only) (Optional) Display entries for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

Required Privilege Level

view

List of Sample Output

show ospf log on page 2131
show ospf log topology voice on page 2131

Output Fields

Table 185 on page 2131 lists the output fields for the **show (ospf | ospf3) log** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2018, Juniper Networks, Inc.

Table 185: show (ospf | ospf3) log Output Fields

Field Name	Field Description
When	Time, in weeks (**w**) and days (**d**), since the SPF calculation was made.
Type	Type of calculation: **Cleanup**, **External**, **Interarea**, **NSSA**, **Redist**, **SPF**, **Stub**, **Total**, or **Virtuallink**.
Elapsed	Amount of time, in seconds, that elapsed during the operation, or the time required to complete the SPF calculation. The start time is the time displayed in the **When** field.

Sample Output

show ospf log

```
user@host> show ospf log
When                Type                Elapsed

1w4d 17:25:58       Stub                0.000017
1w4d 17:25:58       SPF                 0.000070
1w4d 17:25:58       Stub                0.000019
1w4d 17:25:58       Interarea           0.000054
1w4d 17:25:58       External            0.000005
1w4d 17:25:58       Cleanup             0.000203
1w4d 17:25:58         Total             0.000537
1w4d 17:24:48       SPF                 0.000125
1w4d 17:24:48       Stub                0.000017
1w4d 17:24:48       SPF                 0.000100
1w4d 17:24:48       Stub                0.000016
1w4d 17:24:48       Interarea           0.000056
1w4d 17:24:48       External            0.000005
1w4d 17:24:48       Cleanup             0.000238
1w4d 17:24:48         Total             0.000600
...
```

show ospf log topology voice

```
user@host> show ospf log topology voice
Topology voice SPF log:

   Last instance of each event type
When                Type                Elapsed
00:06:11            SPF                 0.000116
00:06:11            Stub                0.000114
00:06:11            Interarea           0.000126
00:06:11            External            0.000067
00:06:11            NSSA                0.000037
00:06:11            Cleanup             0.000186

   Maximum length of each event type
When                Type                Elapsed
00:13:43            SPF                 0.000140
00:13:33            Stub                0.000116
00:13:43            Interarea           0.000128
00:13:33            External            0.000075
```

Copyright © 2018, Juniper Networks, Inc.

```
00:13:38        NSSA            0.000039
00:13:53        Cleanup         0.000657

   Last 100 events
When            Type            Elapsed

00:13:53        SPF             0.000090
00:13:53        Stub            0.000041
00:13:53        Interarea       0.000123
00:13:53        External        0.000040
00:13:53        NSSA            0.000038
00:13:53        Cleanup         0.000657
00:13:53          Total         0.001252
.
.
00:06:11        SPF             0.000116
00:06:11        Stub            0.000114
00:06:11        Interarea       0.000126
00:06:11        External        0.000067
00:06:11        NSSA            0.000037
00:06:11        Cleanup         0.000186
00:06:11          Total         0.000818
```

 Copyright © 2018, Juniper Networks, Inc.

show (ospf | ospf3) neighbor

List of Syntax

Syntax on page 2133
Syntax (EX Series Switches and QFX Series) on page 2133

Syntax

show (ospf | ospf3) neighbor
<brief | detail | extensive>
<area *area-id*>
<instance (all | *instance-name*)>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>
<*neighbor*>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>

Syntax (EX Series Switches and QFX Series)

show (ospf | ospf3) neighbor
<brief | detail | extensive>
<area *area-id*>
<instance (all | *instance-name*)>
<interface *interface-name*>
<*neighbor*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
instance all option introduced in Junos OS Release 9.1.
instance all option introduced in Junos OS Release 9.1 for EX Series switches.
area, **interface**, and **realm** options introduced in Junos OS Release 9.2.
area and **interface** options introduced in Junos OS Release 9.2 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display information about OSPF neighbors.

CPU utilization might increase while the device learns its OSPF neighbors. We recommend that you use the **show (ospf | ospf3) neighbor** command after the device learns and establishes OSPF neighbor adjacencies. Depending on the size of your network, this might take several minutes. If you receive a “timeout communicating with routing daemon” error when using the **show (ospf | ospf3) neighbor** command, wait several minutes before attempting to use the command again. This is not a critical system error, but you might experience a delay in using the CLI.

Options

none—Display standard information about all OSPF neighbors for all routing instances.

brief | detail | extensive—(Optional) Display the specified level of output.

area ***area-id***—(Optional) Display information about the OSPF neighbors for the specified area.

instance (all | ***instance-name*****)**—(Optional) Display all OSPF interfaces for all routing instances or under the named routing instance.

Copyright © 2018, Juniper Networks, Inc.

interface ***interface-name***—(Optional) Display information about OSPF neighbors for the specified logical interface.

logical-system (all | ***logical-system-name*****)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

neighbor—(Optional) Display information about the specified OSPF neighbor.

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)—(OSPFv3 only) (Optional) Display information about the OSPF neighbors for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

Required Privilege Level view

Related Documentation

- clear (ospf | ospf3) neighbor on page 2078

List of Sample Output

show ospf neighbor brief on page 2136
show ospf neighbor detail on page 2136
show ospf neighbor extensive on page 2137
show ospf3 neighbor detail on page 2138
show ospf neighbor area area-id on page 2138
show ospf neighbor interface interface-name on page 2138
show ospf3 neighbor instance all (OSPFv3 Multiple Family Address Support Enabled) on page 2138

Output Fields Table 186 on page 2134 lists the output fields for the **show (ospf | ospf3) neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 186: show (ospf | ospf3) neighbor Output Fields

Field Name	Field Description	Level of Output
Address	Address of the neighbor.	All levels
Interface	Interface through which the neighbor is reachable.	All levels

 Copyright © 2018, Juniper Networks, Inc.

Table 186: show (ospf | ospf3) neighbor Output Fields (continued)

Field Name	Field Description	Level of Output
State	State of the neighbor: • **Attempt**—Valid only for neighbors attached to nonbroadcast networks. It indicates that no recent information has been received from the neighbor, but that a more concerted effort must be made to contact the neighbor. • **Down**—Initial state of a neighbor conversation. It indicates that no recent information has been received from the neighbor. Hello packets might continue to be sent to neighbors in the **Down** state, although at a reduced frequency. • **Exchange**—Routing device is describing its entire link-state database by sending database description packets to the neighbor. Each packet has a sequence number and is explicitly acknowledged. • **ExStart**—First step in creating an adjacency between the two neighboring routing devices. The goal of this step is to determine which routing device is the master, and to determine the initial sequence number. • **Full**—Neighboring routing devices are fully adjacent. These adjacencies appear in router link and network link advertisements. • **Init**—A hello packet has recently been sent by the neighbor. However, bidirectional communication has not yet been established with the neighbor. This state may occur, for example, because the routing device itself did not appear in the neighbor's hello packet. • **Loading**—Link-state request packets are sent to the neighbor to acquire more recent advertisements that have been discovered (but not yet received) in the **Exchange** state. • **2Way**—Communication between the two routing devices is bidirectional. This state has been ensured by the operation of the Hello Protocol. This is the most advanced state short of beginning adjacency establishment. The (backup) designated router is selected from the set of neighbors in state **2Way** or greater.	All levels
ID	Router ID of the neighbor.	All levels
Pri	Priority of the neighbor to become the designated router.	All levels
Dead	Number of seconds until the neighbor becomes unreachable.	All levels
Link state acknowledgment list	Number of link-state acknowledgments received.	**extensive**
Link state retransmission list	Total number of link-state advertisements retransmitted. For **extensive** output only, the following information is also displayed: • **Type**—Type of link advertisement: **ASBR**, **Sum**, **Extern**, **Network**, **NSSA**, **OpaqArea**, **Router**, or **Summary**. • **LSA ID**—LSA identifier included in the advertisement. An asterisk preceding the identifier marks database entries that originated from the local routing device. • **Adv rtr**—Address of the routing device that sent the advertisement. • **Seq**—Link sequence number of the advertisement.	**detail extensive**

Table 186: show (ospf | ospf3) neighbor Output Fields (continued)

Field Name	Field Description	Level of Output
Neighbor-address	(OSPFv3 only) If the neighbor uses virtual links, the **Neighbor-address** is the site-local, local, or global address. If the neighbor uses a physical interface, the **Neighbor-address** is an IPv6 link-local address.	**detail extensive**
area	Area that the neighbor is in.	**detail extensive**
OSPF3-Intf-Index	(OSPFv3 only) Displays the OSPFv3 interface index.	**detail extensive**
opt	Option bits received in the hello packets from the neighbor.	**detail extensive**
DR or **DR-ID**	Address of the designated router.	**detail extensive**
BDR or **BDR-ID**	Address of the backup designated router.	**detail extensive**
Up	Length of time since the neighbor came up.	**detail extensive**
adjacent	Length of time since the adjacency with the neighbor was established.	**detail extensive**
SPRING Adjacency Labels	Segment routing in networking adjacency labels. NOTE: Displayed only when segment routing is enabled	**detail extensive**
Label	Segment routing label.	**detail extensive**
Flags	Segment routing flags. Flags VL indicate *value* and *local*.	**detail extensive**

Sample Output

show ospf neighbor brief

```
user@host> show ospf neighbor brief
  Address          Intf                State     ID              Pri  Dead
192.168.254.225  fxp3.0              2Way      10.250.240.32   128   36
192.168.254.230  fxp3.0              Full      10.250.240.8    128   38
192.168.254.229  fxp3.0              Full      10.250.240.35   128   33
10.1.1.129       fxp2.0              Full      10.250.240.12   128   37
10.1.1.131       fxp2.0              Full      10.250.240.11   128   38
10.1.2.1         fxp1.0              Full      10.250.240.9    128   32
10.1.2.81        fxp0.0              Full      10.250.240.10   128   33
```

show ospf neighbor detail

```
user@host> show ospf neighbor detail
Address          Interface              State     ID               Pri  Dead
10.0.6.60        lt-1/2/0.12            Full      1.1.1.60         128    38
  Area 0.0.0.0, opt 0x52, DR 0.0.0.0, BDR 0.0.0.0
  Up 23:53:47, adjacent 23:53:34
  SPRING Adjacency Labels:

    Label       Flags
```


Copyright © 2018, Juniper Networks, Inc.

```
    299968      VL

10.0.10.70       lt-1/2/0.14           Full      1.1.1.70         128    37
  Area 0.0.0.0, opt 0x52, DR 0.0.0.0, BDR 0.0.0.0
  Up 23:53:47, adjacent 23:53:47
  SPRING Adjacency Labels:

     Label       Flags

     299952      VL
```

show ospf neighbor extensive

```
user@host> show ospf neighbor extensive
  Address          Interface              State      ID               Pri  Dead
10.5.1.2         ge-1/2/0.1            Full      10.5.1.2         128   33
  area 0.0.0.1, opt 0x42, DR 10.5.1.2, BDR 10.5.1.1
  Up 06:09:42, adjacent 05:17:50
    Link state retransmission list:

      Type      LSA ID            Adv rtr          Seq

     Summary  10.8.56.0        172.25.27.82     0x8000004d

     Router   10.5.1.94        10.5.1.94        0x8000005c

     Network  10.5.24.2        10.5.1.94        0x80000036

     Summary  10.8.57.0        172.25.27.82     0x80000024

     Extern   1.10.90.0        10.8.1.2         0x80000041

     Extern   1.4.109.0        10.6.1.2         0x80000041

     Router   10.5.1.190       10.5.1.190       0x8000005f

     Network  10.5.48.2        10.5.1.190       0x8000003d

     Summary  10.8.58.0        172.25.27.82     0x8000004d

     Extern   1.10.91.0        10.8.1.2         0x80000041

     Extern   1.4.110.0        10.6.1.2         0x80000041

     Router   10.5.1.18        10.5.1.18        0x8000005f

     Network  10.5.5.2         10.5.1.18        0x80000033

     Summary  10.8.59.0        172.25.27.82     0x8000003a

     Summary  10.8.62.0        172.25.27.82     0x80000025

10.5.10.2        ge-1/2/0.10           ExStart   10.5.1.38        128   38
  area 0.0.0.1, opt 0x42, DR 10.5.10.2, BDR 10.5.10.1
  Up 06:09:42
    master, seq 0xac1530f8, rexmit DBD in 2 sec
    rexmit LSREQ in 0 sec
10.5.11.2        ge-1/2/0.11           Full      10.5.1.42        128   33
  area 0.0.0.1, opt 0x42, DR 10.5.11.2, BDR 10.5.11.1
  Up 06:09:42, adjacent 05:27:00
```

Copyright © 2018, Juniper Networks, Inc.

```
Link state retransmission list:

  Type      LSA ID           Adv rtr          Seq

 Summary  10.8.58.0        172.25.27.82     0x8000004d

 Extern   1.10.91.0        10.8.1.2         0x80000041

 Extern   1.1.247.0        10.5.1.2         0x8000003f

 Extern   1.4.110.0        10.6.1.2         0x80000041

 Router   10.5.1.18        10.5.1.18        0x8000005f

 Network  10.5.5.2         10.5.1.18        0x80000033

 Summary  10.8.59.0        172.25.27.82     0x8000003a
```

show ospf3 neighbor detail

```
user@host> show ospf3 neighbor detail
ID               Interface              State     Pri   Dead
10.255.71.13     fe-0/0/2.0             Full      128   30
  Neighbor-address fe80::290:69ff:fe9b:e002
  area 0.0.0.0, opt 0x13, OSPF3-Intf-Index 2
  DR-ID 10.255.71.13, BDR-ID 10.255.71.12
  Up 02:51:43, adjacent 02:51:43
```

show ospf neighbor area area-id

```
user@host >show ospf neighbor area 1.1.1.1
Address          Interface              State     ID               Pri  Dead
192.168.37.47    so-0/0/0.0             Full      10.255.245.4     128    33
  Area 1.1.1.1
192.168.37.55    so-1/0/0.0             Full      10.255.245.5     128    37
  Area 1.1.1.1
```

show ospf neighbor interface interface-name

```
user@host >show ospf neighbor interface so-0/0/0.0
Address          Interface              State     ID               Pri  Dead
192.168.37.47    so-0/0/0.0             Full      10.255.245.4     128    37
  Area 0.0.0.0
192.168.37.47    so-0/0/0.0             Full      10.255.245.4     128    33
  Area 1.1.1.1
192.168.37.47    so-0/0/0.0             Full      10.255.245.4     128    32
  Area 2.2.2.2
```

show ospf3 neighbor instance all (OSPFv3 Multiple Family Address Support Enabled)

```
user @host >  show ospf3 neighbor instance all
Instance: ina
    Realm: ipv6-unicast
    ID               Interface              State     Pri   Dead
   100.1.1.1        fe-0/0/2.0             Full      128     37
       Neighbor-address fe80::217:cb00:c87c:8c03
   Instance: inb
   Realm: ipv4-unicast
   ID               Interface              State     Pri   Dead
  100.1.2.1        fe-0/0/2.1             Full      128     33
```

 Copyright © 2018, Juniper Networks, Inc.

```
Neighbor-address fe80::217:cb00:c97c:8c03
```

show (ospf | ospf3) overview

List of Syntax

Syntax on page 2140
Syntax (EX Series Switch and QFX Series) on page 2140

Syntax

show (ospf | ospf3) overview
<brief | extensive>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>

Syntax (EX Series Switch and QFX Series)

show (ospf | ospf3) overview
<brief | extensive>
<instance *instance-name*>

Release Information

Command introduced in Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
realm option introduced in Junos OS Release 9.2.
Database protection introduced in Junos 10.2.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display Open Shortest Path First (OSPF) overview information.

Options

none—Display standard information about all OSPF neighbors for all routing instances.

brief | extensive—(Optional) Display the specified level of output.

instance *instance-name*—(Optional) Display all OSPF interfaces under the named routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)—(Optional) (OSPFv3 only) Display information about the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

Required Privilege Level

view

List of Sample Output

show ospf overview on page 2142
show ospf overview (With Database Protection) on page 2143
show ospf3 overview (With Database Protection) on page 2144
show ospf overview extensive on page 2144

Copyright © 2018, Juniper Networks, Inc.

Output Fields Table 187 on page 2141 lists the output fields for the **show ospf overview** command. Output fields are listed in the approximate order in which they appear.

Table 187: show ospf overview Output Fields

Field name	Field Description	Level of Output
Instance	OSPF routing instance.	All levels
Router ID	Router ID of the routing device.	All levels
Route table index	Route table index.	All levels
Configured overload	Overload capability is enabled. If the overload timer is also configured, display the time that remains before it is set to expire. This field is not displayed after the timer expires.	All levels
Topology	Topology identifier.	All levels
Prefix export count	Number of prefixes exported into OSPF.	All levels
Full SPF runs	Number of complete Shortest Path First calculations.	All levels
SPF delay	Delay before performing consecutive Shortest Path First calculations.	All levels
SPF holddown	Delay before performing additional Shortest Path First (SPF) calculations after the maximum number of consecutive SPF calculations is reached.	All levels
SPF rapid runs	Maximum number of Shortest Path First calculations that can be performed in succession before the hold-down timer begins.	All levels
LSA refresh time	Refresh period for link-state advertisement (in minutes).	All levels
SPRING	Source protocol routing in networking: **enable** or **disable**.	All levels
Node Segments	Nodes of source protocol routing in networking:**enable** or **disable**.	All levels
Ipv4 Index	Ipv4 Index.	All levels
Index Range	Ipv4 Index range.	All levels
Node Segment Blocks Allocated	Details about node segment blocks.	All levels
Database protection state	Current state of database protection.	All levels
Warning threshold	Threshold at which a warning message is logged (percentage of maximum LSA count).	All levels
Non self-generated LSAs	Number of LSAs whose router ID is not equal to the local router ID: **Current**, **Warning** (threshold), and **Allowed**.	All levels

Copyright © 2018, Juniper Networks, Inc.

Table 187: show ospf overview Output Fields (continued)

Field name	Field Description	Level of Output
Ignore time	How long the database has been in the ignore state.	All levels
Reset time	How long the database must stay out of the ignore or isolated state before it returns to normal operations.	All levels
Ignore count	Number of times the database has been in the ignore state: **Current** and **Allowed**.	All levels
Restart	Graceful restart capability: **enabled** or **disabled**.	All levels
Restart duration	Time period for complete reacquisition of OSPF neighbors.	All levels
Restart grace period	Time period for which the neighbors should consider the restarting routing device as part of the topology.	All levels
Graceful restart helper mode	(OSPFv2) Standard graceful restart helper capability (based on RFC 3623): **enabled** or **disabled**.	All levels
Restart-signaling helper mode	(OSPFv2) Restart signaling-based graceful restart helper capability (based on RFC 4811, RFC 4812, and RFC 4813): **enabled** or **disabled**.	All levels
Helper mode	(OSPFv3) Graceful restart helper capability: **enabled** or **disabled**.	All levels
Trace options	OSPF-specific trace options.	**extensive**
Trace file	Name of the file to receive the output of the tracing operation.	**extensive**
Area	Area number. Area 0.0.0.0 is the backbone area.	All levels
Stub type	Stub type of area: **Normal Stub**, **Not Stub**, or **Not so Stubby Stub**.	All levels
Authentication Type	Type of authentication: **None**, **Password**, or **MD5**. NOTE: The **Authentication Type** field refers to the authentication configured at the **[edit protocols ospf area *area-id*]** level. Any authentication configured for an interface in this area will not affect the value of this field.	All levels
Area border routers	Number of area border routers.	All levels
Neighbors	Number of autonomous system boundary routers.	All levels

Sample Output

show ospf overview

```
user@host> show ospf overview
Instance: master
    Router ID: 10.255.245.6
    Route table index: 0
   Configured overload, expires in 118 seconds
```


Copyright © 2018, Juniper Networks, Inc.

```
    LSA refresh time: 50 minutes
    SPRING: Enabled
    Node Segments: Enabled
    Ipv4 Index : 10, Index Range: 2048
    Node Segment Blocks Allocated:
    Start Index : 0, Size : 256, Label-Range: [ 802048, 802303 ]
    Start Index : 256, Size : 256, Label-Range: [ 802304, 802559 ]
    Start Index : 512, Size : 256, Label-Range: [ 802560, 802815 ]
    Start Index : 768, Size : 256, Label-Range: [ 802816, 803071 ]
    Start Index : 1024, Size : 256, Label-Range: [ 803072, 803327 ]
    Start Index : 1280, Size : 256, Label-Range: [ 803328, 803583 ]
    Start Index : 1536, Size : 256, Label-Range: [ 803584, 803839 ]
    Start Index : 1792, Size : 256, Label-Range: [ 803840, 804095 ]
  Restart: Enabled
     Restart duration: 20 sec
     Restart grace period: 40 sec
     Helper mode: enabled
  Area: 0.0.0.0
      Stub type: Not Stub
      Authentication Type: None
      Area border routers: 0, AS boundary routers: 0
      Neighbors
        Up (in full state): 0
    Topology: default (ID 0)
      Prefix export count: 0
      Full SPF runs: 1
      SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
```

show ospf overview (With Database Protection)

```
user@host> show ospf overview
Instance: master
  Router ID: 10.255.112.218
  Route table index: 0
  LSA refresh time: 50 minutes
  Traffic engineering
  Restart: Enabled
    Restart duration: 180 sec
    Restart grace period: 210 sec
    Graceful restart helper mode: Enabled
    Restart-signaling helper mode: Enabled
  Database protection state: Normal
    Warning threshold: 70 percent
    Non self-generated LSAs: Current 582, Warning 700, Allowed 1000
    Ignore time: 30, Reset time: 60
    Ignore count: Current 0, Allowed 1
  Area: 0.0.0.0
    Stub type: Not Stub
    Authentication Type: None
    Area border routers: 0, AS boundary routers: 0
    Neighbors
      Up (in full state): 160
  Topology: default (ID 0)
    Prefix export count: 0
    Full SPF runs: 70
    SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
    Backup SPF: Not Needed
```

Copyright © 2018, Juniper Networks, Inc.

show ospf3 overview (With Database Protection)

```
user@host> show ospf3 overview
Instance: master
  Router ID: 10.255.112.128
  Route table index: 0
  LSA refresh time: 50 minutes
  Database protection state: Normal
    Warning threshold: 80 percent
    Non self-generated LSAs: Current 3, Warning 8, Allowed 10
    Ignore time: 30, Reset time: 60
    Ignore count: Current 0, Allowed 2
  Area: 0.0.0.0
    Stub type: Not Stub
    Area border routers: 0, AS boundary routers: 0
    Neighbors
      Up (in full state): 1
  Topology: default (ID 0)
    Prefix export count: 0
    Full SPF runs: 7
    SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
    Backup SPF: Not Needed
```

show ospf overview extensive

```
user@host> show ospf overview extensive
Instance: master
  Router ID: 1.1.1.103
  Route table index: 0
  Full SPF runs: 13, SPF delay: 0.200000 sec
  LSA refresh time: 50 minutes
  Restart: Disabled
  Trace options: lsa
  Trace file: /var/log/ospf size 131072 files 10
  Area: 0.0.0.0
    Stub type: Not Stub
    Authentication Type: None
    Area border routers: 0, AS boundary routers: 0
    Neighbors
      Up (in full state): 1
```

 Copyright © 2018, Juniper Networks, Inc.

show (ospf | ospf3) route

List of Syntax

Syntax on page 2145
Syntax (EX Series Switch and QFX Series) on page 2145

Syntax

show (ospf | ospf3) route
<brief | detail | extensive>
<abr | asbr | extern | inter | intra>
<*destination*>
<instance (default | ipv4-multicast | *instance-name*)>
<logical-system (default | ipv4-multicast | *logical-system-name*)>
<network>
<no-backup-coverage>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>
<router>
<topology (default | ipv4-multicast | *topology-name*)>
<transit>

Syntax (EX Series Switch and QFX Series)

show (ospf | ospf3) route
<brief | detail | extensive>
<abr | asbr | extern | inter | intra>
<*destination*>
<instance *instance-name*
<network>
<no-backup-coverage>
<router>
<topology (default | ipv4-multicast | *topology-name*)>
<transit>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
topology option introduced in Junos OS Release 9.0.
realm option introduced in Junos OS Release 9.2.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display the entries in the Open Shortest Path First (OSPF) routing table.

Options

none—Display standard information about all entries in the OSPF routing table for all routing instances and all topologies.

destination—Display routes to the specified IP address (with optional destination prefix length).

brief | detail | extensive—(Optional) Display the specified level of output.

abr—(Optional) Display routes to area border routers.

asbr—(Optional) Display routes to autonomous system border routers.

extern—(Optional) Display external routes.

Copyright © 2018, Juniper Networks, Inc.

inter—(Optional) Display interarea routes.

intra—(Optional) Display intra-area routes.

instance (default | ipv4-multicast | *instance-name*)—(Optional) Display entries for the default routing instance, the IPv4 multicast routing instance, or for the specified routing instance.

logical-system (default | ipv4-multicast | *logical-system-name*)—(Optional) Perform this operation on the default logical system, the IPv4 multicast logical system, or on a particular logical system.

network—(Optional) Display routes to networks.

no-backup-coverage—(Optional) Display routes with no backup coverage.

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)—(OSPFv3 only) (Optional) Display entries in the routing table for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

router—(Optional) Display routes to all routers.

topology (default | ipv4-multicast | *topology-name*)—(OSPFv2 only) (Optional) Display routes for the default OSPF topology, IPv4 multicast topology, or for a particular topology.

transit—(Optional) (OSPFv3 only) Display OSPFv3 routes to pseudonodes.

Required Privilege Level

view

List of Sample Output

show ospf route on page 2148
show ospf route detail on page 2149
show ospf route extensive on page 2149
show ospf3 route on page 2149
show ospf3 route detail on page 2150
show ospf route topology voice on page 2150

Output Fields

Table 188 on page 2146 list the output fields for the **show (ospf | ospf3) route** command. Output fields are listed in the approximate order in which they appear.

Table 188: show (ospf | ospf3) route Output Fields

Field Name	Field Description	Output Level
Topology	Name of the topology.	All levels
Prefix	Destination of the route.	All levels

Copyright © 2018, Juniper Networks, Inc.

Table 188: show (ospf | ospf3) route Output Fields (continued)

Field Name	Field Description	Output Level
Path type	How the route was learned: • **Inter**—Interarea route • **Ext1**—External type 1 route • **Ext2**—External type 2 route • **Intra**—Intra-area route	All levels
Route type	The type of routing device from which the route was learned: • **AS BR**—Route to AS border router. • **Area BR**—Route to area border router. • **Area/AS BR**—Route to router that is both an **Area BR** and **AS BR**. • **Network**—Network router. • **Router**—Route to a router that is neither an **Area BR** nor an **AS BR**. • **Transit**—(OSPFv3 only) Route to a pseudonode representing a transit network, LAN, or nonbroadcast multiaccess (NBMA) link. • **Discard**—Route to a summary discard.	All levels
NH Type	Next-hop type: **LSP** or **IP**.	All levels
Metric	Route's metric value.	All levels
NH-interface	(OSPFv3 only) Interface through which the route's next hop is reachable.	All levels
NH-addr	(OSPFv3 only) IPv6 address of the next hop.	All levels
NextHop Interface	(OSPFv2 only) Interface through which the route's next hop is reachable.	All levels
Nexthop addr/label	(OSPFv2 only) If the **NH Type** is **IP**, then it is the address of the next hop. If the **NH Type** is **LSP**, then it is the name of the label-switched path.	All levels
Area	Area ID of the route.	**detail**
Origin	Router from which the route was learned.	**detail**
Type 7	Route was learned through a not-so-stubby area (NSSA) link-state advertisement (LSA).	**detail**
P-bit	Route was learned through NSSA LSA and the propagate bit was set.	**detail**
Fwd NZ	Forwarding address is nonzero. **Fwd NZ** is only displayed if the route is learned through an NSSA LSA.	**detail**

Table 188: show (ospf | ospf3) route Output Fields (continued)

Field Name	Field Description	Output Level
optional-capability	Optional capabilities propagated in the router LSA. This field is in the output for intra-area router routes only (when **Route Type** is **Area BR**, **AS BR**, **Area/AS BR**, or **Router**), not for interarea router routes or network routes. Three bits in this field are defined as follows: • **0x4 (V)**—Routing device is at the end of a virtual active link. • **0x2 (E)**—Routing device is an autonomous system boundary router. • **0x1 (B)**—Routing device is an area border router.	**detail**
priority	The priority assigned to the prefix: • **high** • **medium** • **low** NOTE: The **priority** field applies only to routes of type **Network**.	**detail**
BGP-ORR Generation-ID	Display the BGP-ORR generation identifier of the main OSPF route. This field is shown only for non-zero values.	**extensive**

Sample Output

show ospf route

```
user@host> show ospf route
Topology default Route Table:

Prefix               Path  Route      NH        Metric NextHop       Nexthop
                     Type  Type       Type             Interface     Address/LSP
1.1.1.60/32          Intra Network    Spring         6 lt-1/2/0.14   10.0.10.70
                                      Bkup SPRING      lt-1/2/0.12   10.0.6.60
1.1.1.70/32          Intra Network    IP             1 lt-1/2/0.14   10.0.10.70
                                      Bkup LSP                       (null)
1.1.1.70/32          Intra Network    Spring         1 lt-1/2/0.14   10.0.10.70
                                      Bkup SPRING      lt-1/2/0.12   10.0.6.60
1.1.1.80/32          Intra Network    IP             6 lt-1/2/0.14   10.0.10.70
                                      Bkup IP          lt-1/2/0.12   10.0.6.60
1.1.1.80/32          Intra Network    Spring         6 lt-1/2/0.14   10.0.10.70

802068 (S=0)         Intra Network    Mpls           0 lt-1/2/0.14   10.0.10.70
                                      Bkup MPLS        lt-1/2/0.12   10.0.6.60
802078 (S=0)         Intra Network    Mpls           0 lt-1/2/0.14   10.0.10.70
                                      Bkup MPLS        lt-1/2/0.12   10.0.6.60
802088 (S=0)         Intra Network    Mpls           0 lt-1/2/0.14   10.0.10.70
                                      Bkup MPLS        lt-1/2/0.12   10.0.6.60
802098 (S=0)         Intra Network    Mpls           0 lt-1/2/0.14   10.0.10.70
                                      Bkup MPLS        lt-1/2/0.12   10.0.6.60
802108 (S=0)         Intra Network    Mpls           0 lt-1/2/0.14   10.0.10.70
                                      Bkup MPLS        lt-1/2/0.12   10.0.6.60
802118 (S=0)         Intra Network    Mpls           0 lt-1/2/0.14   10.0.10.70
                                      Bkup MPLS        lt-1/2/0.12   10.0.6.60
802118               Intra Network    Mpls           0 lt-1/2/0.14   10.0.10.70
                                      Bkup MPLS        lt-1/2/0.12   10.0.6.60
```

 Copyright © 2018, Juniper Networks, Inc.

show ospf route detail

```
user@host> show ospf route detail
Topology default Route Table:

Prefix               Path   Route        NH   Metric  NextHop       Nexthop
                     Type   Type         Type         Interface     addr/label
10.255.14.174        Inter  AS BR        IP      210  t1-3/0/1.0
   area 0.0.0.2, origin 10.255.14.185
10.255.14.178        Intra  Router       IP      200  t3-3/1/3.0
   area 0.0.0.2, origin 10.255.14.178, optional-capability 0x0
10.210.1.0/30        Intra  Network      IP       10  t3-3/1/2.0
   area 0.0.0.2, origin 10.255.14.172, priority medium
198.51.100/24         Inter  Network      IP      210  t1-3/0/1.0
   area 0.0.0.2, origin 10.255.14.185, priority low
192.0.2.0/24         Ext2   Network      IP        0  t1-3/0/1.0
   area 0.0.0.0, origin 10.255.14.174, priority high
203.3.113.0/24        Inter  Network      IP      220  t1-3/0/1.0
   area 0.0.0.2, origin 10.255.14.185, priority high
```

show ospf route extensive

```
user@host> show ospf route extensive
Topology default Route Table:

Prefix               Path  Route       NH       Metric NextHop        Nexthop
                     Type  Type        Type            Interface      Address/LSP
1.1.1.1              Intra Router      IP          100 ge-0/0/2.0     10.1.1.1
  area 0.0.0.0, origin 1.1.1.1, optional-capability 0x0
1.1.1.1/32           Intra Network     IP          100 ge-0/0/2.0     10.1.1.1
  area 0.0.0.0, origin 1.1.1.1, priority medium
BGP-ORR generation-id: 1
```

show ospf3 route

```
user@host> show ospf3 route
Prefix               Path   Route       NH   Metric  NextHop       Nexthop

                     Type   Type        Type          Interface     addr/label
10.255.71.13         Intra  Router      IP        1
  NH-interface fe-0/0/2.0, NH-addr fe80::290:69ff:fe9b:e002
10.255.71.13;0.0.0.2
10.255.245.1         Intra  Router      IP       40  fxp1.1        192.168.36.17

  area 0.0.0.0, origin 10.255.245.1 optional-capability 0x0,
10.255.245.3         Intra  AS BR       IP        1  fxp2.3        192.168.36.34

  area 0.0.0.0, origin 10.255.245.3 optional-capability 0x0,
10.255.245.1/32      Intra  Network     IP       40  fxp1.1        192.168.36.17

  area 0.0.0.0, origin 10.255.245.1, priority high
10.255.245.2/32      Intra  Network     IP        0  lo0.0
  area 0.0.0.0, origin 10.255.245.2, priority medium
10.255.245.3/32      Intra  Network     IP        1  fxp2.3        192.168.36.34

  area 0.0.0.0, origin 10245.3, priority low
                     Intra  Transit     IP        1
  NH-interface fe-0/0/2.0
192::168:222:84/126  Intra  Network     IP        1
  NH-interface fe-0/0/2.0
```

```
abcd::71:12/128       Intra  Network     IP       0
  NH-interface lo0.0
abcd::71:13/128       Intra  Network     LSP      1
  NH-interface fe-0/0/2.0, NH-addr lsp-cd
```

show ospf3 route detail

```
user@host> show ospf3 route detail
Prefix                                       Path   Route        NH    Metric
                                             type   type         type
10.255.14.174                                Intra  Area/AS BR   IP    110
   NH-interface so-1/2/2.0
   Area 0.0.0.0, Origin 10.255.14.174, Optional-capability 0x3
10.255.14.178                                Intra  Router       IP    200
   NH-interface t3-3/1/3.0
   Area 0.0.0.0, Origin 10.255.14.178, Optional-capability 0x0
10.255.14.185;0.0.0.2                        Intra  Transit      IP    200
   NH-interface t1-3/0/1.0
   NH-interface so-1/2/2.0
   Area 0.0.0.0, Origin 10.255.14.185
1000:1:1::1/128                              Inter  Network      IP    110
   NH-interface so-1/2/2.0
   Area 0.0.0.0, Origin 10.255.14.174, Priority low
1001:2:1::/48                                Ext1   Network      IP    110
   NH-interface so-1/2/2.0
   Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority medium
1002:1:7::/48                                Ext2   Network      IP    0
   NH-interface so-1/2/2.0
   Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority low
1002:3:4::/48                                Ext2   Network      IP    0
   NH-interface so-1/2/2.0
   Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority high
abcd::10:255:14:172/128                      Intra  Network      IP    0
   NH-interface lo0.0
   Area 0.0.0.0, Origin 10.255.14.172, Priority low
```

show ospf route topology voice

```
user@host show ospf route topology voice
Topology voice Route Table:
Prefix             Path   Route       NH   Metric  NextHop       Nexthop
                   Type   Type        Type         Interface     addr/label
10.255.8.2         Intra  Router      IP        1  so-0/2/0.0
10.255.8.3         Intra  Router      IP        2  so-0/2/0.0
10.255.8.1/32      Intra  Network     IP        0  lo0.0
10.255.8.2/32      Intra  Network     IP        1  so-0/2/0.0
10.255.8.3/32      Intra  Network     IP        2  so-0/2/0.0
192.168.8.0/29     Intra  Network     IP        2  so-0/2/0.0
192.168.8.44/30    Intra  Network     IP        2  so-0/2/0.0
192.168.8.46/32    Intra  Network     IP        1  so-0/2/0.0
192.168.8.48/30    Intra  Network     IP        1  so-0/2/1.0
192.168.8.52/30    Intra  Network     IP        2  so-0/2/0.0
192.168.9.44/30    Intra  Network     IP        1  so-0/2/0.0
192.168.9.45/32    Intra  Network     IP        2  so-0/2/0.0
```

 Copyright © 2018, Juniper Networks, Inc.

show (ospf | ospf3) statistics

List of Syntax

Syntax on page 2151
Syntax (EX Series Switch and QFX Series) on page 2151

Syntax

show (ospf | ospf3) statistics
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>

Syntax (EX Series Switch and QFX Series)

show (ospf | ospf3) statistics
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
realm option introduced in Junos OS Release 9.2.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display OSPF statistics.

Options

none—Display OSPF statistics for all routing instances.

instance *instance-name*—(Optional) Display all statistics for the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)—(Optional) (OSPFv3 only) Display all statistics for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

Required Privilege Level

view

Related Documentation

- clear (ospf | ospf3) statistics on page 2081

List of Sample Output

show ospf statistics on page 2153
show ospf statistics logical-system all on page 2153
show ospf3 statistics on page 2154

Output Fields

Table 189 on page 2152 lists the output fields for the **show (ospf | ospf3) statistics** command. Output fields are listed in the approximate order in which they appear.

Table 189: show (ospf | ospf3) statistics Output Fields

Field Name	Field Description
Packet type	Type of OSPF packet.
Total Sent/Total Received	Total number of packets sent and received.
Last 5 seconds Sent/Last 5 seconds Received	Total number of packets sent and received in the last 5 seconds.
DBDs retransmitted	Total number of database description packets retransmitted, and number retransmitted in the last 5 seconds.
LSAs flooded	Total number of link-state advertisements flooded, and number flooded in the last 5 seconds.
LSAs flooded high-prio	Total number of high priority link-state advertisements flooded, and number flooded in the last 5 seconds. A link-state advertisement is deemed a high priority if it has changed since it was last sent.
LSAs retransmitted	Total number of link-state advertisements retransmitted, and number retransmitted in the last 5 seconds.
LSAs transmitted to nbr	Total number of link-state advertisements transmitted to a neighbor, and number transmitted in the last 5 seconds.
LSAs requested	Total number of link-state advertisements requested by neighboring devices, and number requested in the last 5 seconds.
LSAs acknowledged	Total number of link-state advertisements acknowledged, and number acknowledged in the last 5 seconds.
Flood queue depth	Total number of entries in the extended queue.
Total rexmit entries	Total number of retransmission entries waiting to be sent from the OSPF routing instance.
db summaries	Total number of database description summaries waiting to be sent from the OSPF routing instance.
lsreq entries	Total number of link-state request entries waiting to be sent from the OSPF routing instance.

 Copyright © 2018, Juniper Networks, Inc.

Table 189: show (ospf | ospf3) statistics Output Fields (continued)

Field Name	Field Description
Receive errors	Number and type of receive errors. Some sample receive errors include: • **mtu mismatches** • **no interface found** • **no virtual link found** • **nssa mismatches** • **stub area mismatches** • **subnet mismatches** If there are no receive errors, the output displays **none**.

Sample Output

show ospf statistics

```
user@host> show ospf statistics
Packet type             Total                 Last 5 seconds
                   Sent       Received         Sent       Received
   Hello             31             14            2              2
     DbD              9             10            0              0
   LSReq              2              2            0              0
LSUpdate              8             16            0              0
   LSAck              9              9            0              0

DBDs retransmitted     :                    3, last 5 seconds :          0
LSAs flooded           :                   12, last 5 seconds :          0
LSAs flooded high-prio :                    0, last 5 seconds :          0
LSAs retransmitted     :                    0, last 5 seconds :          0
LSAs transmitted to nbr:                    3, last 5 seconds :          0
LSAs requested         :                    5, last 5 seconds :          0
LSAs acknowledged      :                   19, last 5 seconds :          0

Flood queue depth      :               0
Total rexmit entries   :               0
db summaries           :               0
lsreq entries          :               0

Receive errors:
  862 no interface found
  115923 no virtual link found
```

show ospf statistics logical-system all

```
user@host> show ospf statistics logical-system all
logical-system: C
OSPF instance is not running
-----

logical-system: B

Packet type             Total                 Last 5 seconds
                   Sent       Received         Sent       Received
```

Copyright © 2018, Juniper Networks, Inc.

```
   Hello        313740       313653           1             0
     DbD             3            2           0             0
   LSReq             1            1           0             0
LSUpdate          2752         1825           0             0
   LSAck          1821         2747           0             0

DBDs retransmitted     :                  0, last 5 seconds :          0
LSAs flooded           :               2741, last 5 seconds :          0
LSAs flooded high-prio :                 10, last 5 seconds :          0
LSAs retransmitted     :                  0, last 5 seconds :          0
LSAs transmitted to nbr:                  2, last 5 seconds :          0
LSAs requested         :                  1, last 5 seconds :          0
LSAs acknowledged      :               1831, last 5 seconds :          0

Flood queue depth      :               0
Total rexmit entries   :               0
db summaries           :               0
lsreq entries          :               0

Receive errors:
  None
-----

logical-system: A

Packet type             Total                   Last 5 seconds
                   Sent       Received         Sent       Received
   Hello        313698       313695           0             0
     DbD             2            3           0             0
   LSReq             1            1           0             0
LSUpdate          1825         2752           0             0
   LSAck          2747         1821           0             0

DBDs retransmitted     :                  0, last 5 seconds :          0
LSAs flooded           :               1825, last 5 seconds :          0
LSAs flooded high-prio :                 10, last 5 seconds :          0
LSAs retransmitted     :                  0, last 5 seconds :          0
LSAs transmitted to nbr:                  1, last 5 seconds :          0
LSAs requested         :                  2, last 5 seconds :          0
LSAs acknowledged      :               2748, last 5 seconds :          0

Flood queue depth      :               0
Total rexmit entries   :               0
db summaries           :               0
lsreq entries          :               0

Receive errors:
  None
-----
```

show ospf3 statistics

```
user@host> show ospf3 statistics
Packet type             Total                   Last 5 seconds
                   Sent       Received         Sent       Received
   Hello             0            0           0             0
     DbD             0            0           0             0
   LSReq             0            0           0             0
LSUpdate             0            0           0             0
   LSAck             0            0           0             0
```

 Copyright © 2018, Juniper Networks, Inc.

```
DBDs retransmitted     :                    0, last 5 seconds :          0
LSAs flooded           :                    0, last 5 seconds :          0
LSAs flooded high-prio :                    0, last 5 seconds :          0
LSAs retransmitted     :                    0, last 5 seconds :          0
LSAs transmitted to nbr:                    0, last 5 seconds :          0
LSAs requested         :                    0, last 5 seconds :          0
LSAs acknowledged      :                    0, last 5 seconds :          0

Flood queue depth      :               0
Total rexmit entries   :               0
db summaries           :               0
lsreq entries          :               0

Receive errors:
  None
```

 Copyright © 2018, Juniper Networks, Inc.

CHAPTER 25

Protocol-Independent Routing Operational Commands

- show route label
- show route label-switched-path
- show route localization
- show route martians
- show route match-prefix
- show route next-hop
- show route no-community
- show route output
- show route protocol
- show route range
- show route receive-protocol
- show route resolution
- show route snooping
- show route source-gateway
- show route summary
- show route table
- show route terse

Copyright © 2018, Juniper Networks, Inc.

show as-path

List of Syntax

Syntax on page 2159
Syntax (EX Series Switches) on page 2159

Syntax

show as-path
<brief | detail>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show as-path
<brief | detail>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display the distribution of autonomous system (AS) paths that the local routing device is using (usually through the routing table). Use this command to debug problems for AS paths and to understand how AS paths have been manipulated through a policy (through the **as-path-prepend** action) or through aggregation.

AS paths are stored in a hash table. A hash table is one method for fast lookup. Each entry in the table is called a bucket. Junos OS computes a hash value that indicates in which bucket the AS path is stored. The AS paths are dispersed among the hash buckets so that a manageable number of AS paths is stored in each bucket. Only unique AS paths are stored. Duplicate AS paths increase a reference count, but do not increase the number of AS paths stored in the hash table.

Options

none—Display basic information about AS paths that the local routing device is using (same as brief).

brief | detail—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- show as-path summary on page 2167

List of Sample Output

show as-path on page 2160
show as-path detail on page 2161

Output Fields

Table 190 on page 2160 lists the output fields for the **show as-path** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2018, Juniper Networks, Inc.

Table 190: show as-path Output Fields

Field Name	Field Description	Level of Output
Total AS paths	Total number of AS paths.	**brief** none
Bucket	Bucket number.	All levels
Count	Number of AS path entries in this bucket.	All levels
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated. • **Atomic**—Route is an aggregate of several route prefixes. • **Aggregator**—Routing device has summarized a range of prefixes.	All levels
domain	Number of independent AS domains. The AS paths of an independent AS domain are not shared with the AS paths and AS path attributes of other domains, including the master routing instance domain.	**detail**
neighbor as	AS peer address.	**detail**
length	Length of the AS path.	**detail**
segments	Length of the AS segment descriptor.	**detail**
unique-count	Number of unique autonomous systems (ASs) present in the AS path	**detail**
references	Path reference count.	**detail**

Sample Output

show as-path

```
user@host> show as-path
Total AS paths: 30382
  Bucket 0    Count: 36
    I
    14203 2914 174 31752 I
    14203 2914 701 21512 I
    14203 2914 1239 26632 I
    14203 2914 1239 29704 I
    14203 2914 4323 10248 I
    14203 2914 4766 23560 I
    14203 2914 6395 32776 I
    14203 2914 7911 11272 I
    14203 2914 12180 18440 I
    14203 2914 17408 17416 I
    14203 2914 701 702 24586 I
    14203 2914 1239 4657 9226 I
```


Copyright © 2018, Juniper Networks, Inc.

```
    14203 2914 1239 7132 16394 I
    14203 2914 1299 8308 34826 I
    14203 2914 3320 5603 28682 I
    14203 2914 3491 1680 33802 I
    14203 2914 3549 7908 27658 I
    14203 2914 3549 20804 30730 I
    14203 2914 7018 2687 9226 I
    14203 2914 174 9318 9318 23564 I
    14203 2914 701 3786 3786 23564 I
    14203 2914 701 4761 4795 9228 I
    14203 2914 1239 7132 5673 18444 I
    14203 2914 3491 20485 24588 24588 I
    14203 2914 5511 2200 1945 2060 I
    14203 2914 7911 14325 14325 14348 I
    14203 2914 701 4637 9230 9230 9230 I
    14203 2914 6395 14 14 14 14 I
    14203 2914 9299 6163 6163 6163 6163 9232 I
    14203 2914 3356 3356 3356 3356 3356 11955 21522 I
    14203 2914 9837 9837 9219 I Aggregator: 9219 202.27.91.253
    14203 2914 174 30209 30222 30222 30222 ?
    14203 2914 1299 5377 I (Atomic) Aggregator: 5377 193.219.192.22
    14203 2914 4323 36097 I (Atomic) Aggregator: 36097 216.69.252.254
    14203 2914 209 2516 17676 23813 I (Atomic) Aggregator: 23813 219.127.233.66
  Bucket 1    Count: 28
    14203 2914 35847 I
    14203 2914 174 19465 I
    14203 2914 174 35849 I
    14203 2914 2828 32777 I
    14203 2914 4323 14345 I
    14203 2914 4323 29705 I
    14203 2914 6395 32777 I

...
```

show as-path detail

```
user@host> show as-path detail
Total AS paths: 30410
  Bucket 0    Count: 36
    AS path: I
      domain 0, length 0, segments 0, unique-count 0, references 54
    AS path: 14203 2914 174 31752 I
      domain 1, neighbor as: 14203, length 4, segments 1, unique-count 4,
references 2
    AS path: 14203 2914 701 21512 I
      domain 1, neighbor as: 14203, length 4, segments 1, unique-count 4,
references 2
    AS path: 14203 2914 1239 26632 I
      domain 1, neighbor as: 14203, length 4, segments 1, unique-count 5,
references 2
    AS path: 14203 2914 1239 29704 I
      domain 1, neighbor as: 14203, length 4, segments 1, unique-count 4,
references 2
    AS path: 14203 2914 4323 10248 I
      domain 1, neighbor as: 14203, length 4, segments 1, unique-count 6,
references 2
    AS path: 14203 2914 4766 23560 I
      domain 1, neighbor as: 14203, length 4, segments 1, unique-count 4,
references 2
    AS path: 14203 2914 6395 32776 I
      domain 1, neighbor as: 14203, length 4, segments 1, unique-count 5,
```

```
references 3
    AS path: 14203 2914 7911 11272 I
      domain 1, neighbor as: 14203, length 4, segments 1, unique-count 6,
references 2
    AS path: 14203 2914 12180 18440 I
      domain 1, neighbor as: 14203, length 4, segments 1, unique-count 3,
references 3
    AS path: 14203 2914 17408 17416 I
      domain 1, neighbor as: 14203, length 4, segments 1, unique-count 8,
references 3
    AS path: 14203 2914 701 702 24586 I
      domain 1, neighbor as: 14203, length 5, segments 1, unique-count 4,
references 3
    AS path: 14203 2914 1239 4657 9226 I
      domain 1, neighbor as: 14203, length 5, segments 1, unique-count 5,
references 7
    AS path: 14203 2914 1239 7132 16394 I
      domain 1, neighbor as: 14203, length 5, segments 1, unique-count 7,
references 2
    AS path: 14203 2914 1299 8308 34826 I
      domain 1, neighbor as: 14203, length 5, segments 1, unique-count 8,
references 2
    AS path: 14203 2914 3320 5603 28682 I
      domain 1, neighbor as: 14203, length 5, segments 1, unique-count 4,
references 2
    AS path: 14203 2914 3491 1680 33802 I
      domain 1, neighbor as: 14203, length 5, segments 1, unique-count 14,
references 2
    AS path: 14203 2914 3549 7908 27658 I
      domain 1, neighbor as: 14203, length 5, segments 1, unique-count 6,
references 2
    AS path: 14203 2914 3549 20804 30730 I
      domain 1, neighbor as: 14203, length 5, segments 1, unique-count 24,
references 2
    AS path: 14203 2914 7018 2687 9226 I
      domain 1, neighbor as: 14203, length 5, segments 1, unique-count 4,
references 3
    AS path: 14203 2914 174 9318 9318 23564 I
      domain 1, neighbor as: 14203, length 6, segments 1, unique-count 4,
references 2
    AS path: 14203 2914 701 3786 3786 23564 I
      domain 1, neighbor as: 14203, length 6, segments 1, unique-count 4,
references 2
    AS path: 14203 2914 701 4761 4795 9228 I
      domain 1, neighbor as: 14203, length 6, segments 1, unique-count 4,
references 14
    AS path: 14203 2914 1239 7132 5673 18444 I
      domain 1, neighbor as: 14203, length 6, segments 1, unique-count 4,
references 2
    AS path: 14203 2914 3491 20485 24588 24588 I
      domain 1, neighbor as: 14203, length 6, segments 1, unique-count 4,
references 4
    AS path: 14203 2914 5511 2200 1945 2060 I
      domain 1, neighbor as: 14203, length 6, segments 1, unique-count 4,
references 2
    AS path: 14203 2914 7911 14325 14325 14348 I
      domain 1, neighbor as: 14203, length 6, segments 1, unique-count 4,
references 2
    AS path: 14203 2914 701 4637 9230 9230 9230 I
      domain 1, neighbor as: 14203, length 7, segments 1, unique-count 4,
references 3
```

 Copyright © 2018, Juniper Networks, Inc.

```
    AS path: 14203 2914 6395 14 14 14 14 I
      domain 1, neighbor as: 14203, length 7, segments 1, unique-count 4,
references 10
...
```

show as-path domain

List of Syntax
Syntax on page 2164
Syntax (EX Series Switches) on page 2164

Syntax

show as-path domain
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show as-path domain

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display autonomous system (AS) path domain information.

Options

none—(Optional) Display AS path domain information for all routing instances.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show as-path domain on page 2166

Output Fields

Table 191 on page 2164 lists the output fields for the **show as-path domain** command. Output fields are listed in the approximate order in which they appear

Table 191: show as-path domain Output Fields

Field Name	Field Description
Domain	Number of independent AS domains. The AS paths of an independent AS domain are not shared with the AS paths and AS path attributes of other domains, including the master routing instance domain.
Primary	Primary AS number.
References	Path reference count.
***Number* Paths**	Number of known AS paths.

Copyright © 2018, Juniper Networks, Inc.

Table 191: show as-path domain Output Fields (continued)

Field Name	Field Description
Flags	Information about the AS path: • **ASLoop**—Path contains an AS loop. • **Atomic**—Path includes the ATOMIC_AGGREGATE path attribute. • **Local**—Path was created by local aggregation. • **Master**—Path was created by the master routing instance.
Local AS	AS number of the local routing device.
Loops	How many times this AS number can appear in an AS path.

Sample Output

show as-path domain

```
user@host> show as-path domain
Domain: 1        Primary: 10458
    References:          3 Paths:      30383
    Flags: Master
    Local AS: 10458  Loops: 1
```


Copyright © 2018, Juniper Networks, Inc.

show as-path summary

List of Syntax

Syntax on page 2167
Syntax (EX Series Switches) on page 2167

Syntax

show as-path summary
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show as-path summary

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display autonomous system (AS) path summary information.

AS paths are stored in a hash table. A hash table is one method for fast lookup. Each entry in the table is called a bucket. Junos OS computes a hash value that indicates in which bucket the AS path is stored. The AS paths are dispersed among the hash buckets so that a manageable number of AS paths is stored in each bucket. Only unique AS paths are stored. Duplicate AS paths increase a reference count, but do not increase the number of AS paths stored in the hash table.

Options

none—(Optional) Display AS path summary information for all routing instances.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- show as-path on page 2159

List of Sample Output

show as-path summary on page 2168

Output Fields

Table 192 on page 2167 lists the output fields for the **show as-path summary** command. Output fields are listed in the approximate order in which they appear.

Table 192: show as-path summary Output Fields

Field Name	Field Description
AS Paths	Number of AS paths.
Buckets	Number of hash buckets in use.
Max	Maximum number of AS path entries per bucket.

Table 192: show as-path summary Output Fields (continued)

Field Name	Field Description
Min	Minimum number of AS path entries per bucket.
Avg	Average number of AS path entries per bucket.
Std deviation	Standard deviation of AS path entries per bucket.

Sample Output

show as-path summary

```
user@host> show as-path summary
AS Paths  Buckets      Max     Min     Avg   Std deviation
   30425     1024       95      12      29        6.481419
```

 Copyright © 2018, Juniper Networks, Inc.

show route

List of Syntax

Syntax on page 2169
Syntax (EX Series Switches) on page 2169

Syntax

show route
<all>
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>
<private>
<te-ipv4-prefix-ip *te-ipv4-prefix-ip*>
<te-ipv4-prefix-node-ip *te-ipv4-prefix-node-ip*>
<te-ipv4-prefix-node-iso *te-ipv4-prefix-node-iso*>

Syntax (EX Series Switches)

show route
<all>
<*destination-prefix*>
<private>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Option **private** introduced in Junos OS Release 9.5.
Option **private** introduced in Junos OS Release 9.5 for EX Series switches.
Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.
Option **display-client-data** introduced in Junos OS Release 16.2R1 on MX80, MX104, MX240, MX480, MX960, MX2010, MX2020, vMX Series routers.
Options **te-ipv4-prefix-ip**, **te-ipv4-prefix-node-ip**, and **te-ipv4-prefix-node-iso** introduced in Junos OS Release 17.2R1 on MX Series and PTX Series.

Description

Display the active entries in the routing tables.

Options

none—Display brief information about all active entries in the routing tables.

all—(Optional) Display information about all routing tables, including private, or internal, routing tables.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

private—(Optional) Display information only about all private, or internal, routing tables.

display-client-data —(Optional) Display client id and cookie information for routes installed by the routing protocol process client applications.

te-ipv4-prefix-ip *te-ipv4-prefix-ip*—(Optional) Display IPv4 address of the traffic-engineering prefix, without the mask length if present in the routing table.

Copyright © 2018, Juniper Networks, Inc.

te-ipv4-prefix-node-ip ***te-ipv4-prefix-node-ip***—(Optional) Display all prefixes that have originated from the traffic-engineering node. You can filter IPv4 node addresses from the traffic-engineered routes in the **lsdist.0** table.

te-ipv4-prefix-node-iso ***te-ipv4-prefix-node-iso***—(Optional) Display all prefixes that have originated from the traffic-engineering node. You can filter IPv4 routes with the specified ISO circuit ID from the **lsdist.0** table.

Required Privilege Level

view

Related Documentation

- *Understanding IS-IS Configuration*
- *Example: Configuring IS-IS*
- *Examples: Configuring Internal BGP Peering*
- *Examples: Configuring External BGP Peering*
- *Examples: Configuring OSPF Routing Policy*
- *Verifying and Managing Junos OS Enhanced Subscriber Management*

List of Sample Output

Output Fields

Table 193 on page 2170 describes the output fields for the **show route** command. Output fields are listed in the approximate order in which they appear.

Table 193: show route Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).
***number* destinations**	Number of destinations for which there are routes in the routing table.

 Copyright © 2018, Juniper Networks, Inc.

Table 193: show route Output Fields (continued)

Field Name	Field Description
number **routes**	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active). • **holddown** (routes that are in the pending state before being declared inactive). A holddown route was once the active route and is no longer the active route. The route is in the holddown state because a protocol still has interest in the route, meaning that the interest bit is set. A protocol might have its interest bit set on the previously active route because the protocol is still advertising the route. The route will be deleted after all protocols withdraw their advertisement of the route and remove their interest bit. A persistent holddown state often means that the interested protocol is not releasing its interest bit properly. However, if you have configured advertisement of multiple routes (with the **add-path** or **advertise-inactive** statement), the holddown bit is most likely set because BGP is advertising the route as an active route. In this case, you can ignore the holddown state because nothing is wrong. • **hidden** (routes that are not used because of a routing policy).

destination-prefix	Route destination (for example:10.0.0.1/24). Sometimes the route information is presented in another format, such as:
	• *MPLS-label* (for example, 80001).
	• *interface-name* (for example, ge-1/0/2).
	• *neighbor-address*:*control-word-status*:*encapsulation type*:*vc-id* :*source* (Layer 2 circuit only. For example, 10.1.1.195:NoCtrlWord:1:1:Local/96):
	  ◦ *neighbor-address*—Address of the neighbor.
	  ◦ *control-word-status*—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**.
	  ◦ *encapsulation type*—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.
	  ◦ *vc-id*—Virtual circuit identifier.
	  ◦ *source*—Source of the advertisement: **Local** or **Remote**.

[*protocol*, *preference* **]**	Protocol from which the route was learned and the preference value for the route. • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.
weeks:days hours:minutes:seconds	How long the route been known (for example, **2w4d 13:11:14**, or 2 weeks, 4 days, 13 hours, 11 minutes, and 14 seconds).
metric	Cost value of the indicated route. For routes within an AS, the cost is determined by the IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.

Table 193: show route Output Fields (continued)

<table>
<tr><th>Field Name</th><th>Field Description</th></tr>
<tr><td>localpref</td><td>Local preference value included in the route.</td></tr>
<tr><td>from</td><td>Interface from which the route was received.</td></tr>
<tr><td>AS path</td><td>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:
• I—IGP.
• E—EGP.
• ?—Incomplete; typically, the AS path was aggregated.
When AS path numbers are included in the route, the format is as follows:
• [ ]—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured.
• { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.
• ( )—Parentheses enclose a confederation.
• ([ ])—Parentheses and brackets enclose a confederation set.
NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</td></tr>
<tr><td>encapsulated</td><td>Extended next-hop encoding capability enabled for the specified BGP community for routing IPv4 traffic over IPv6 tunnels. When BGP receives routes without the tunnel community, IPv4-0ver IPv6 tunnels are not created and BGP routes are resolved without encapsulation.</td></tr>
<tr><td>Route Labels</td><td>Stack of labels carried in the BGP route update.</td></tr>
<tr><td>validation-state</td><td>(BGP-learned routes) Validation status of the route:
• Invalid—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database.
• Unknown—Indicates that the prefix is not among the prefixes or prefix ranges in the database.
• Unverified—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers.
• Valid—Indicates that the prefix and autonomous system pair are found in the database.</td></tr>
<tr><td>to</td><td>Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.
If the destination is Discard, traffic is dropped.</td></tr>
</table>

 Copyright © 2018, Juniper Networks, Inc.

Table 193: show route Output Fields (continued)

Field Name	Field Description
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word **Selected**. This field can also contain the following information: • **Weight**—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • **Balance**—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing. • ***lsp-path-name***—Name of the LSP used to reach the next hop. • ***label-action***—MPLS label and operation occurring at the next hop. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label). For VPNs, expect to see multiple **push** operations, corresponding to the inner and outer labels required for VPN routes (in the case of a direct PE-to-PE connection, the VPN route would have the inner label push only).
Private unicast	(Enhanced subscriber management for MX Series routers) Indicates that an access-internal route is managed by enhanced subscriber management. By contrast, access-internal routes *not* managed by enhanced subscriber management are displayed with associated next-hop and media access control (MAC) address information.
balance	Distribution of the load based on the underlying operational interface bandwidth for equal-cost multipaths (ECMP) across the nexthop gateways in percentages.

Sample Output

show route

```
user@host> show route
inet.0: 11 destinations, 12 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:65500:1:10.0.0.20/240
                   *[MVPN/70] 19:53:41, metric2 1
                      Indirect
1:65500:1:10.0.0.40/240
                   *[BGP/170] 19:53:29, localpref 100, from 10.0.0.30
                      AS path: I
                    > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
                    [BGP/170] 19:53:26, localpref 100, from 10.0.0.33
                      AS path: I
                    > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
1:65500:1:10.0.0.60/240
                   *[BGP/170] 19:53:29, localpref 100, from 10.0.0.30
                      AS path: I
                    > to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF
                    [BGP/170] 19:53:25, localpref 100, from 10.0.0.33
                      AS path: I
                    > to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF
```

Copyright © 2018, Juniper Networks, Inc.

show route (VPN)

The following sample output shows a VPN route with composite next hops enabled. The first **Push** operation corresponds to the outer label. The second **Push** operation corresponds to the inner label.

user@host> **show route 192.0.2.0**

```
13979:665001.inet.0: 871 destinations, 3556 routes (871 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.0/24        [BGP/170] 00:28:32, localpref 100, from 10.9.9.160
                      AS path: 13980 ?, validation-state: unverified
                    > to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
                    [BGP/170] 00:28:28, localpref 100, from 10.9.9.169
                      AS path: 13980 ?, validation-state: unverified
                    > to 10.100.0.42 via ae2.0, Push 126016, Push 300368(top)
                   #[Multipath/255] 00:28:28, metric2 102
                    > to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
                      to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
```

show route (with Destination Prefix)

user@host> **show route 192.168.0.0/12**

```
inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.0/12      *[Static/5] 2w4d 12:54:27
                    > to 192.168.167.254 via fxp0.0
```

show route destination-prefix detail

user@host> **show route 198.51.100.0 detail**

```
inet.0: 15 destinations, 20 routes (15 active, 0 holddown, 0 hidden)
198.51.100.0/24 (2 entries, 2 announced)
        *BGP    Preference: 170/-101
         ...
         BGP-Static Preference: 4294967292
               Next hop type: Discard
                 Address: 0x9041ae4
                 Next-hop reference count: 2
                 State: <NoReadvrt Int Ext AlwaysFlash>
        Inactive reason: Route Preference
        Local AS:   200
        Age: 4d 1:40:40
        Validation State: unverified
        Task: RT
        Announcement bits (1): 2-BGP_RT_Background
        AS path: 4 5 6 I
```

show route extensive

user@host> **show route extensive**

```
v1.mvpn.0: 5 destinations, 8 routes (5 active, 1 holddown, 0 hidden)
1:65500:1:10.0.0.40/240 (1 entry, 1 announced)
    *BGP   Preference: 170/-101
```


Copyright © 2018, Juniper Networks, Inc.

```
           PMSI: Flags 0x0: Label[0:0:0]: PIM-SM: Sender 10.0.0.40 Group
203.0.113.1
           Next hop type: Indirect
           Address: 0x92455b8
           Next-hop reference count: 2
           Source: 10.0.0.30
           Protocol next hop: 10.0.0.40
           Indirect next hop: 2 no-forward
           State: <Active Int Ext>
                     Local AS: 64510 Peer AS: 64511
           Age: 3  Metric2: 1
          Validation State: unverified
           Task: BGP_64510.10.0.0.30+179
           Announcement bits (2): 0-PIM.v1 1-mvpn global task
           AS path: I (Originator) Cluster list:  10.0.0.30
           AS path:  Originator ID: 10.0.0.40
           Communities: target:64502:100 encapsulation:0L:14           Import
Accepted
           Localpref: 100
           Router ID: 10.0.0.30
           Primary Routing Table bgp.mvpn.0
           Indirect next hops: 1
                Protocol next hop: 10.0.0.40 Metric: 1
                Indirect next hop: 2 no-forward
                Indirect path forwarding next hops: 1
                      Next hop type: Router
                      Next hop: 10.0.24.4 via lt-0/3/0.24 weight 0x1
                10.0.0.40/32 Originating RIB: inet.3
                      Metric: 1                       Node path count: 1
                      Forwarding nexthops: 1
                         Nexthop: 10.0.24.4 via lt-0/3/0.24
```

show route extensive (ECMP)

```
user@host> show route extensive
*IS-IS Preference: 15
                Level: 1
                Next hop type: Router, Next hop index: 1048577
                Address: 0xXXXXXXXXXX
                Next-hop reference count: YY
                Next hop: 198.51.100.2 via ae1.0 balance 43%, selected
                Session Id: 0x141
                Next hop: 192.0.2.2 via ae0.0 balance 57%
```

show route extensive (Multipath Resolution)

```
user@host> show route extensive
inet.0: 37 destinations, 37 routes (36 active, 0 holddown, 1 hidden)
10.1.1.2/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.1.1.2/32 -> {indirect(1048574)}
        *Static Preference: 5
                Next hop type: Indirect, Next hop index: 0
                Address: 0xb39d1b0
                Next-hop reference count: 2
                Next hop type: Router, Next hop index: 581
                Next hop: 10.1.1.2 via ge-2/0/1.0, selected
                Session Id: 0x144
                Next hop: 10.2.1.2 via ge-2/0/2.0, selected
                Session Id: 0x145
```

Copyright © 2018, Juniper Networks, Inc.

```
                Protocol next hop: 10.1.1.1
                Indirect next hop: 0xb2b20f0 1048574 INH Session ID: 0x143
                State: <Active Int Ext>
                Age: 2:53  Metric2: 0
                Validation State: unverified
                Task: RT
                Announcement bits (2): 0-KRT 2-Resolve tree 1
                AS path: I
                Indirect next hops: 1
                        Protocol next hop: 10.1.1.1
                      Indirect next hop: 0xb2b20f0 1048574 INH Session ID: 0x143

                        Indirect path forwarding next hops: 2
                                Next hop type: Router
                                Next hop: 10.1.1.2 via ge-2/0/1.0
                                Session Id: 0x144
                                Next hop: 10.2.1.2 via ge-2/0/2.0
                                Session Id: 0x145
                     10.1.1.1/32 Originating RIB: inet.0
                        Node path count: 1
                          Node flags: 1
                        Forwarding nexthops: 2 (Merged)
                        Nexthop: 10.1.1.2 via ge-2/0/1.0
                         Nexthop: 10.2.1.2 via ge-2/0/2.0
```

```
user@host> show route active-path extensive
user@host> show route 198.51.100.1 active-path extensive

inet.0: 1000061 destinations, 1000082 routes (1000061 active, 0 holddown, 0 hidden)
198.51.100.1/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 198.51.100.1/32 -> {indirect(1051215)}
unicast reverse-path: 0
[ae0.0 ae1.0]
Page 0 idx 0, (group Internet-IPv4 type External) Type 1 val 0xbb2e53d8 (adv_entry)
Advertised metrics:
Nexthop: Self
AS path: [500] 410 I
Communities:
Path 198.51.100.1 from 10.0.0.11 Vector len 4. Val: 0
*BGP Preference: 170/-101
Next hop type: Indirect, Next hop index: 0
Address: 0x2e9aacdc
Next-hop reference count: 500000
Source: 10.0.0.11
Next hop type: Router, Next hop index: 0
Next hop: 10.0.12.2 via ae0.0 weight 0x1
Label operation: Push 3851, Push 25, Push 20(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 25: None; Label 20: None;
Label element ptr: 0xb5dc1780
Label parent element ptr: 0x18d48080
Label element references: 2
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 10.0.12.2 via ae0.0 weight 0x1
Label operation: Push 3851, Push 25, Push 22(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 25: None; Label 22: None;
Label element ptr: 0xb5dc1700
Label parent element ptr: 0x18d41000
```

 Copyright © 2018, Juniper Networks, Inc.

```
Label element references: 2
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 10.0.12.2 via ae0.0 weight 0x1
Label operation: Push 3851, Push 24, Push 48(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 24: None; Label 48: None;
Label element ptr: 0x18d40800
Label parent element ptr: 0x18d49780
Label element references: 3
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 10.0.12.2 via ae0.0 weight 0x1
Label operation: Push 3851, Push 24, Push 49(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 24: None; Label 49: None;
Label element ptr: 0xb5dc1680
Label parent element ptr: 0x18d48f00
Label element references: 2
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 10.0.13.3 via ae1.0 weight 0x1
Label operation: Push 3851, Push 25, Push 21(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 25: None; Label 21: None;
Label element ptr: 0xb5dc1600
Label parent element ptr: 0x18d44d80
Label element references: 2
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 10.0.13.3 via ae1.0 weight 0x1
Label operation: Push 3851, Push 25, Push 25(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 25: None; Label 25: None;
Label element ptr: 0xb5dc1580
Label parent element ptr: 0x18d3da80
Label element references: 2
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 10.0.13.3 via ae1.0 weight 0x1, selected
Label operation: Push 3851, Push 24, Push 68(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 24: None; Label 68: None;
Label element ptr: 0x18d41500
Label parent element ptr: 0x18d49000
Label element references: 3
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 10.0.13.3 via ae1.0 weight 0x1
Label operation: Push 3851, Push 24, Push 69(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 24: None; Label 69: None;
Label element ptr: 0xb5dc1500
Label parent element ptr: 0x18d48300
Label element references: 2
```

Copyright © 2018, Juniper Networks, Inc.

```
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Protocol next hop: 10.0.0.11
Label operation: Push 3851
Label TTL action: prop-ttl
Load balance label: Label 3851: None;
Indirect next hop: 0x1883e200 1051215 INH Session ID: 0xb0d
State:
Local AS: 500 Peer AS: 500
Age: 1:40:03 Metric2: 2
Validation State: unverified
Task: BGP_500.10.0.0.11
Announcement bits (5): 0-KRT 8-KRT 9-BGP_RT_Background 10-Resolve tree 5 11-Resolve
 tree 8
AS path: 410 I
Accepted
Route Label: 3851
Localpref: 100
Router ID: 10.0.0.11
Indirect next hops: 1
Protocol next hop: 10.0.0.11 Metric: 2
Label operation: Push 3851
Label TTL action: prop-ttl
Load balance label: Label 3851: None;
Indirect next hop: 0x1883e200 1051215 INH Session ID: 0xb0d
Indirect path forwarding next hops (Merged): 8
Next hop type: Router
Next hop: 10.0.12.2 via ae0.0 weight 0x1
Session Id: 0x0
Next hop: 10.0.12.2 via ae0.0 weight 0x1
Session Id: 0x0
Next hop: 10.0.12.2 via ae0.0 weight 0x1
Session Id: 0x0
Next hop: 10.0.12.2 via ae0.0 weight 0x1
Session Id: 0x0
Next hop: 10.0.13.3 via ae1.0 weight 0x1
Session Id: 0x0
Next hop: 10.0.13.3 via ae1.0 weight 0x1
Session Id: 0x0
Next hop: 10.0.13.3 via ae1.0 weight 0x1
Session Id: 0x0
Next hop: 10.0.13.3 via ae1.0 weight 0x1
Session Id: 0x0
10.0.0.11/32 Originating RIB: inet.3
Metric: 1 Node path count: 4
Node flags: 1
Indirect nexthops: 4
Protocol Nexthop: 10.0.0.4 Metric: 1 Push 24
Indirect nexthop: 0x1880f200 1048597 INH Session ID: 0xb0c
Path forwarding nexthops link: 0x36120400
Path inh link: 0x0
Indirect path forwarding nexthops: 2
Nexthop: 10.0.12.2 via ae0.0
Session Id: 0
Nexthop: 10.0.13.3 via ae1.0
Session Id: 0
10.0.0.4/32 Originating RIB: inet.3
Metric: 1 Node path count: 1
Forwarding nexthops: 2
Nexthop: 10.0.12.2 via ae0.0
```

 Copyright © 2018, Juniper Networks, Inc.

```
Session Id: 0
Nexthop: 10.0.13.3 via ae1.0
Session Id: 0
Protocol Nexthop: 10.0.0.5 Metric: 1 Push 24
Indirect nexthop: 0x18810000 1048596 INH Session ID: 0xb0b
Path forwarding nexthops link: 0x1545be00
Path inh link: 0x0
Indirect path forwarding nexthops: 2
Nexthop: 10.0.12.2 via ae0.0
Session Id: 0
Nexthop: 10.0.13.3 via ae1.0
Session Id: 0
10.0.0.5/32 Originating RIB: inet.3
Metric: 1 Node path count: 1
Forwarding nexthops: 2
Nexthop: 10.0.12.2 via ae0.0
Session Id: 0
Nexthop: 10.0.13.3 via ae1.0
Session Id: 0
Protocol Nexthop: 10.0.0.6 Metric: 1 Push 25
Indirect nexthop: 0x1880e600 1048588 INH Session ID: 0xb0a
Path forwarding nexthops link: 0x3611f440
Path inh link: 0x0
Indirect path forwarding nexthops: 2
Nexthop: 10.0.12.2 via ae0.0
Session Id: 0
Nexthop: 10.0.13.3 via ae1.0
Session Id: 0
10.0.0.6/32 Originating RIB: inet.3
Metric: 1 Node path count: 1
Forwarding nexthops: 2
Nexthop: 10.0.12.2 via ae0.0
Session Id: 0
Nexthop: 10.0.13.3 via ae1.0
Session Id: 0
Protocol Nexthop: 10.0.0.7 Metric: 1 Push 25
Indirect nexthop: 0x1880dc00 1048586 INH Session ID: 0xb09
Path forwarding nexthops link: 0x15466d80
Path inh link: 0x0
Indirect path forwarding nexthops: 2
Nexthop: 10.0.12.2 via ae0.0
Session Id: 0
Nexthop: 10.0.13.3 via ae1.0
Session Id: 0
10.0.0.7/32 Originating RIB: inet.3
Metric: 1 Node path count: 1
Forwarding nexthops: 2
Nexthop: 10.0.12.2 via ae0.0
Session Id: 0
Nexthop: 10.0.13.3 via ae1.0
Session Id: 0
```

show route (Enhanced Subscriber Management)

```
user@host> show route
inet.0: 41 destinations, 41 routes (40 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

198.51.100.11/24    *[Access-internal/12] 00:00:08
                > to #0 10.0.0.1.93.65 via demux0.1073741824
```

```
198.51.100.12/24    *[Access-internal/12] 00:00:08
                    Private unicast
```

show route (IPv6 Flow Specification)

```
user@host> show route
inet6.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::10:255:185:19/128
                   *[Direct/0] 05:11:27
                    > via lo0.0
2001:db8::11:11:11:0/120
                   *[BGP/170] 00:28:58, localpref 100
                      AS path: 2000 I, validation-state: unverified
                    > to 2001:db8::13:14:2:2 via ge-1/1/4.0
2001:db8::13:14:2:0/120*[Direct/0] 00:45:07
                    > via ge-1/1/4.0
2001:db8::13:14:2:1/128*[Local/0] 00:45:18
                      Local via ge-1/1/4.0
fe80::2a0:a50f:fc71:71d5/128
                   *[Direct/0] 05:11:27
                    > via lo0.0
fe80::5e5e:abff:feb0:933e/128
                   *[Local/0] 00:45:18
                      Local via ge-1/1/4.0

inet6flow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::11:11:11:10/128,*,proto=6,dstport=80,srcport=65535/term:1
                   *[BGP/170] 00:28:58, localpref 100, from 2001:db8::13:14:2:2
                      AS path: 2000 I, validation-state: unverified
                      Fictitious
2001:db8::11:11:11:30/128,*,icmp6-type=128,len=100,dscp=10/term:2
                   *[BGP/170] 00:20:54, localpref 100, from 2001:db8::13:14:2:2
                      AS path: 2000 I, validation-state: unverified
                      Fictitious
```

show route display-client-data detail

```
user@host> show route 198.51.100.0/24 display-client-data detail
inet.0: 59 destinations, 70 routes (59 active, 0 holddown, 0 hidden)
198.51.100.0/24 (1 entry, 1 announced)
        State: <FlashAll>
        *BGP-Static Preference: 5/-101
                Next hop type: Indirect, Next hop index: 0
                Address: 0xa5c2af8
                Next-hop reference count: 2
                Next hop type: Router, Next hop index: 1641
                Next hop: 192.0.2.1 via ge-2/1/1.0, selected
                Session Id: 0x160
                Protocol next hop: 192.0.2.1
                Indirect next hop: 0xa732cb0 1048621 INH Session ID: 0x17e
                State: <Active Int Ext AlwaysFlash NSR-incapable Programmed>
                Age: 3:13       Metric2: 0
                Validation State: unverified
                Announcement bits (3): 0-KRT 5-LDP 6-Resolve tree 3
```

 Copyright © 2018, Juniper Networks, Inc.

```
                AS path: I
                Client id: 1, Cookie: 1
```

show route te-ipv4-prefix-ip

```
user@host> show route te-ipv4-prefix-ip 10.10.10.10
lsdist.0: 283 destinations, 283 routes (283 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both


PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L1:0
 }/1152
                   *[IS-IS/15] 00:01:01
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0101.0101.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                   *[IS-IS/18] 00:01:01
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0202.0202.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                   *[IS-IS/18] 00:01:01
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0303.0303.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                   *[IS-IS/18] 00:01:01
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0404.0404.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                   *[IS-IS/18] 00:01:01
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0505.0505.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                   *[IS-IS/18] 00:01:01
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0606.0606.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                   *[IS-IS/18] 00:01:01
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0707.0707.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                   *[IS-IS/18] 00:01:01
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                   *[IS-IS/18] 00:01:01
                      Fictitious
```

show route te-ipv4-prefix-ip extensive

```
user@host>show route te-ipv4-prefix-ip 10.10.10.10 extensive
lsdist.0: 298 destinations, 298 routes (298 active, 0 holddown, 0 hidden)
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State:<Active NotInstall>
                Local AS:   64496
                Age: 7:58
```

Copyright © 2018, Juniper Networks, Inc.

```
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0x40, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0101.0101.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
       State: <Active NotInstall
       Local AS:   64496
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0>

PREFIX { Node { AS:64496 ISO:0100.0202.0202.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
       Local AS:   64496
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0303.0303.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0404.0404.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
```


Copyright © 2018, Juniper Networks, Inc.

```
                State: <Active NotInstall>
                Local AS:   64496
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0505.0505.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0606.0606.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0707.0707.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
```

Copyright © 2018, Juniper Networks, Inc.

```
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0x40, Algo: 0
```

show route te-ipv4-prefix-node-iso

```
user@host> show route te-ipv4-prefix-node-iso 0100.0a0a.0a0a.00
lsdist.0: 283 destinations, 283 routes (283 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L1:0
 }/1152
                   *[IS-IS/15] 00:05:20
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.1.1.1/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.2.2.2/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.3.3.3/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.4.4.4/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.5.5.5/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.6.6.6/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.7.7.7/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
```

show route te-ipv4-prefix-node-iso extensive

```
user@host> show route te-ipv4-prefix-node-iso 0100.0a0a.0a0a.00 extensive
lsdist.0: 283 destinations, 283 routes (283 active, 0 holddown, 0 hidden)
PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L1:0
```

 Copyright © 2018, Juniper Networks, Inc.

```
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0x40, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.1.1.1/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1001, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.2.2.2/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1002, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.3.3.3/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
```

```
                Prefix SID: 1003, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.4.4.4/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1004, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.5.5.5/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1005, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.6.6.6/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1006, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.7.7.7/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:47
```

 Copyright © 2018, Juniper Networks, Inc.

```
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1007, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0x40, Algo: 0
```

show route te-ipv4-prefix-node-iso detail

```
user@host> show route te-ipv4-prefix-node-iso 0100.0a0a.0a0a.00 detail
lsdist.0: 283 destinations, 283 routes (283 active, 0 holddown, 0 hidden)
PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L1:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0x40, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.1.1.1/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1001, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.2.2.2/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
```

```
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1002, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.3.3.3/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1003, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.4.4.4/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1004, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.5.5.5/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1005, Flags: 0xe0, Algo: 0
```

 Copyright © 2018, Juniper Networks, Inc.

```
PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.6.6.6/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1006, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.7.7.7/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1007, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:64496 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   64496
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0x40, Algo: 0
```

show route active-path

List of Syntax

Syntax on page 2190
Syntax (EX Series Switches) on page 2190

Syntax

show route active-path
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route active-path
<brief | detail | extensive | terse>

Release Information

Command introduced in Junos OS Release 8.0.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display all active routes for destinations. An active route is a route that is selected as the best path. Inactive routes are not displayed.

Options

none—Display all active routes.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show route active-path on page 2190
show route active-path brief on page 2191
show route active-path detail on page 2191
show route active-path extensive on page 2192
show route active-path terse on page 2194

Output Fields

For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route active-path

```
user@host> show route active-path

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.70.19/32    *[Direct/0] 21:33:52
```

 Copyright © 2018, Juniper Networks, Inc.

```
                        > via lo0.0
10.255.71.50/32     *[IS-IS/15] 00:18:13, metric 10
                        > to 172.16.100.1 via so-2/1/3.0
172.16.100.1/24       *[Direct/0] 00:18:36
                        > via so-2/1/3.0
172.16.100.1/32       *[Local/0] 00:18:41
                          Local via so-2/1/3.0
192.168.64.0/21     *[Direct/0] 21:33:52
                        > via fxp0.0
192.168.70.19/32    *[Local/0] 21:33:52
                          Local via fxp0.0
```

show route active-path brief

The output for the **show route active-path brief** command is identical to that for the **show route active-path** command. For sample output, see show route active-path on page 2190.

show route active-path detail

```
user@host> show route active-path detail

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)

10.255.70.19/32 (1 entry, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:   200
                Age: 21:37:10
                Task: IF
                Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

                AS path: I

10.255.71.50/32 (1 entry, 1 announced)
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Router, Next hop index: 397
                Next-hop reference count: 4
                Next hop: 172.16.100.1 via so-2/1/3.0, selected
                State: <Active Int>
                Local AS:   200
                Age: 21:31      Metric: 10
                Task: IS-IS
                Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
                AS path: I

172.16.100.0/24 (1 entry, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via so-2/1/3.0, selected
                State: <Active Int>
                Local AS:   200
                Age: 21:54
                Task: IF
                Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
```

Copyright © 2018, Juniper Networks, Inc.

```
                AS path: I

172.16.100.1/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 11
                Interface: so-2/1/3.0
                State: <Active NoReadvrt Int>
                Local AS:   200
                Age: 21:59
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

192.168.64.0/21 (1 entry, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via fxp0.0, selected
                State: <Active Int>
                Local AS:   200
                Age: 21:37:10
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

192.168.70.19/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 11
                Interface: fxp0.0
                State: <Active NoReadvrt Int>
                Local AS:   200
                Age: 21:37:10
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I
```

show route active-path extensive

```
user@host> show route active-path extensive

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
10.255.70.19/32 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:   200
                Age: 21:39:47
                Task: IF
                Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

                AS path: I

10.255.71.50/32 (1 entry, 1 announced)
```

 Copyright © 2018, Juniper Networks, Inc.

```
TSI:
KRT in-kernel 10.255.71.50/32 -> {172.16.100.1}
IS-IS level 2, LSP fragment 0
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Router, Next hop index: 397
                Next-hop reference count: 4
                Next hop: 172.16.100.1 via so-2/1/3.0, selected
                State: <Active Int>
                Local AS:   200
                Age: 24:08      Metric: 10
                Task: IS-IS
                Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
                AS path: I

172.16.100.1/24 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via so-2/1/3.0, selected
                State: <Active Int>
                Local AS:   200
                Age: 24:31
                Task: IF
                Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

                AS path: I

172.16.100.1/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 11
                Interface: so-2/1/3.0
                State: <Active NoReadvrt Int>
                Local AS:   200
                Age: 24:36
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

192.168.64.0/21 (1 entry, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via fxp0.0, selected
                State: <Active Int>
                Local AS:   200
                Age: 21:39:47
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

192.168.70.19/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 11
                Interface: fxp0.0
```

```
State: ‹Active NoReadvrt Int›
Local AS:   200
Age: 21:39:47
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I
```

show route active-path terse

user@host> **show route active-path terse**

```
inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.255.70.19/32    D   0                       >lo0.0
* 10.255.71.50/32    I  15         10            >172.16.100.1.
* 172.16.100.0/24       D    0                      >so-2/1/3.0
* 172.16.100.2/32       L    0                       Local
* 192.168.64.0/21    D   0                       >fxp0.0
* 192.168.70.19/32   L   0                        Local
```


Copyright © 2018, Juniper Networks, Inc.

show route advertising-protocol

Syntax

show route advertising-protocol *protocol neighbor-address*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display the routing information as it has been prepared for advertisement to a particular neighbor of a particular dynamic routing protocol.

Options

brief | detail | extensive | terse—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

neighbor-address—Address of the neighboring router to which the route entry is being transmitted.

protocol—Protocol transmitting the route:

- **bgp**—Border Gateway Protocol
- **dvmrp**—Distance Vector Multicast Routing Protocol
- **msdp**—Multicast Source Discovery Protocol
- **pim**—Protocol Independent Multicast
- **rip**—Routing Information Protocol
- **ripng**—Routing Information Protocol next generation

Additional Information

Routes displayed are routes that the routing table has exported into the routing protocol and that have been filtered by the associated protocol's **export** routing policy statements. Starting with Junos OS Release 13.3, you can display the routing instance table **foo** for any address family, on a VPN route reflector, or a VPN AS boundary router that is advertising local VPN routes. However, If you do not specify the **table** in the command, the output displays each VRF prefix twice.

Required Privilege Level

view

Related Documentation

- *Example: Configuring the MED Attribute That Determines the Exit Point in an AS*

List of Sample Output

show route advertising-protocol bgp (Layer 3 VPN) on page 2198
show route advertising-protocol bgp detail on page 2198

Copyright © 2018, Juniper Networks, Inc.

show route advertising-protocol bgp detail (Aggregate Extended Community Bandwidth) on page 2198
show route advertising-protocol bgp detail (Labeled Unicast) on page 2199
show route advertising-protocol bgp detail (Layer 2 VPN) on page 2199
show route advertising-protocol bgp detail (Layer 3 VPN) on page 2199
show route advertising-protocol bgp extensive all (Next Hop Self with RIB-out IP Address) on page 2200

Output Fields Table 194 on page 2196 lists the output fields for the **show route advertising-protocol** command. Output fields are listed in the approximate order in which they appear.

Table 194: show route advertising-protocol Output Fields

Field Name	Field Description	Level of Output
routing-table-name	Name of the routing table—for example, inet.0.	All levels
number destinations	Number of destinations for which there are routes in the routing table.	All levels
number routes	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active) • **holddown** (routes that are in the pending state before being declared inactive) • **hidden** (routes that are not used because of a routing policy)	All levels
Prefix	Destination prefix.	**brief** none
destination-prefix **(entry , announced)**	Destination prefix. The **entry** value is the number of routes for this destination, and the **announced** value is the number of routes being announced for this destination.	**detail extensive**
BGP group and type	BGP group name and type (**Internal** or **External**).	**detail extensive**
Route Distinguisher	Unique 64-bit prefix augmenting each IP subnet.	**detail extensive**
Advertised Label	Incoming label advertised by the Label Distribution Protocol (LDP). When an IP packet enters a label-switched path (LSP), the ingress router examines the packet and assigns it a label based on its destination, placing the label in the packet's header. The label transforms the packet from one that is forwarded based on its IP routing information to one that is forwarded based on information associated with the label.	**detail extensive**
Label-Base, range	First label in a block of labels and label block size. A remote PE router uses this first label when sending traffic toward the advertising PE router.	**detail extensive**
VPN Label	Virtual private network (VPN) label. Packets are sent between CE and PE routers by advertising VPN labels. VPN labels transit over either a Resource Reservation Protocol (RSVP) or a Label Distribution Protocol (LDP) label-switched path (LSP) tunnel.	**detail extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 194: show route advertising-protocol Output Fields (continued)

Field Name	Field Description	Level of Output
Nexthop	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route. If the next-hop advertisement to the peer is **Self**, and the RIB-out next hop is a specific IP address, the RIB-out IP address is included in the extensive output. See show route advertising-protocol bgp extensive all (Next Hop Self with RIB-out IP Address) on page 2200.	All levels
MED	Multiple exit discriminator value included in the route.	**brief**
Lclpref or **Localpref**	Local preference value included in the route.	All levels
Queued	When BGP route prioritization is enabled and a route is present in a priority queue, this shows which priority queue the route is in.	All levels except **brief**
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the local AS number associated with the AS path if configured on the router, or if AS path prepending is configured. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.	All levels
Route Labels	Stack of labels carried in the BGP route update.	**detail extensive**
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.	**detail extensive**
Originator ID	(For route reflected output only) Address of routing device that originally sent the route to the route reflector.	**detail extensive**
Communities	Community path attribute for the route. See the output field table for the show route detail command for all possible values for this field.	**detail extensive**
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.	**detail extensive**

Table 194: show route advertising-protocol Output Fields (continued)

Field Name	Field Description	Level of Output
Attrset AS	Number, local preference, and path of the autonomous system (AS) that originated the route. These values are stored in the **Attrset** attribute at the originating router.	**detail extensive**
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).	**detail extensive**
control flags	Control flags: **none** or **Site Down**.	**detail extensive**
mtu	Maximum transmission unit (MTU) of the Layer 2 circuit.	**detail extensive**

Sample Output

show route advertising-protocol bgp (Layer 3 VPN)

```
user@host> show route advertising-protocol bgp 10.255.14.171
      VPN-A.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
      Prefix             Nexthop                 MED    Lclpref AS path
      10.255.14.172/32   Self                      1        100 I
      VPN-B.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
      Prefix             Nexthop                 MED    Lclpref AS path
      10.255.14.181/32   Self                      2        100 I
```

show route advertising-protocol bgp detail

```
user@host> show route advertising-protocol bgp 111.222.1.3 detail
bgp20.inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
111.222.1.11/32 (1 entry, 1 announced)
 BGP group pe-pe type Internal
     Route Distinguisher: 111.255.14.11:69
     Advertised Label: 100000
     next hop: Self
     Localpref: 100
     AS path: 2 I
     Communities: target:69:20
     AIGP 210
111.8.0.0/16 (1 entry, 1 announced)
 BGP group pe-pe type Internal
     Route Distinguisher: 111.255.14.11:69
     Advertised Label: 100000
     Next hop: Self
     Localpref: 100
     AS path: 2 I
     Communities: target:69:20
     AIGP 210
```

show route advertising-protocol bgp detail (Aggregate Extended Community Bandwidth)

```
user@host> show route advertising-protocol bgp 10.0.4.2 10.0.2.0/30 detail
inet.0: 20 destinations, 26 routes (20 active, 0 holddown, 0 hidden)
* 10.0.2.0/30 (2 entries, 1 announced)
 BGP group external2 type External
     Nexthop: Self
```

 Copyright © 2018, Juniper Networks, Inc.

```
     AS path: [65000] 65001 I
     Communities: bandwidth:65000:80000000
```

show route advertising-protocol bgp detail (Labeled Unicast)

```
user@host>show route advertising bgp 1.1.1.3 detail
inet.0: 69 destinations, 70 routes (69 active, 0 holddown, 0 hidden)
* 1.1.1.8/32 (2 entries, 2 announced)
BGP group ibgp type Internal
Route Labels: 1000123(top) 1000124 1000125 1000126
Nexthop: 1.1.1.4
MED: 7
Localpref: 100
AS path: [5] I
Cluster ID: 3.3.3.3
Originator ID: 1.1.1.1
Entropy label capable
inet6.0: 26 destinations, 28 routes (26 active, 0 holddown, 0 hidden)
* 100::1/128 (2 entries, 1 announced)
BGP group ibgp type Internal
Labels: 1000123(top) 1000124 1000125 1000126
Nexthop: ::ffff:1.1.1.4
Localpref: 100
AS path: [5] I
Cluster ID: 3.3.3.3
Originator ID: 1.1.1.1
```

show route advertising-protocol bgp detail (Layer 2 VPN)

```
user@host> show route advertising-protocol bgp 192.168.24.1 detail
vpn-a.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
192.168.16.1:1:1:1/96 (1 entry, 1 announced)
 BGP group int type Internal
     Route Distinguisher: 192.168.16.1:1
     Label-base : 32768, range : 3
     Nexthop: Self
     Localpref: 100
     AS path: I
     Communities: target:65412:100
     AIGP 210
     Layer2-info: encaps:VLAN, control flags:, mtu:
```

show route advertising-protocol bgp detail (Layer 3 VPN)

```
user@host> show route advertising-protocol bgp 10.255.14.176 detail
vpna.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
* 10.49.0.0/30 (1 entry, 1 announced)
 BGP group ibgp type Internal
     Route Distinguisher: 10.255.14.174:2
     VPN Label: 101264
     Nexthop: Self
     Localpref: 100
     AS path: I
     Communities: target:200:100
     AIGP 210
     AttrSet AS: 100
         Localpref: 100
         AS path: I
...
```

show route advertising-protocol bgp extensive all (Next Hop Self with RIB-out IP Address)

```
user@host> show route advertising-protocol bgp 200.0.0.2 170.0.1.0/24 extensive all
inet.0: 13 destinations, 19 routes (13 active, 0 holddown, 6 hidden)
  170.0.1.0/24 (2 entries, 1 announced)
 BGP group eBGP-INTEROP type External
     Nexthop: Self (rib-out 10.100.3.2)
     AS path: [4713] 200 I
...
```

 Copyright © 2018, Juniper Networks, Inc.

show route all

List of Syntax

Syntax on page 2201
Syntax (EX Series Switches) on page 2201

Syntax

show route all
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route all

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display information about all routes in all routing tables, including private, or internal, tables.

Options

none—Display information about all routes in all routing tables, including private, or internal, tables.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- show route brief on page 2208
- show route detail on page 2220

List of Sample Output

show route all on page 2201

Output Fields

In Junos OS Release 9.5 and later, only the output fields for the **show route all** command display all routing tables, including private, or hidden, routing tables. The output field table of the show route command does not display entries for private, or hidden, routing tables in Junos OS Release 9.5 and later.

Sample Output

show route all

The following example displays a snippet of output from the **show route** command and then displays the same snippet of output from the **show route all** command:

```
user@host> show route
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
0                  *[MPLS/0] 2d 02:24:39, metric 1
```

```
                      Receive
1                  *[MPLS/0] 2d 02:24:39, metric 1
                      Receive
2                  *[MPLS/0] 2d 02:24:39, metric 1
                      Receive
800017             *[VPLS/7] 1d 14:00:16
                    > via vt-3/2/0.32769, Pop
800018             *[VPLS/7] 1d 14:00:26
                    > via vt-3/2/0.32772, Pop
```

```
user@host> show route all
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
0                  *[MPLS/0] 2d 02:19:12, metric 1
                      Receive
1                  *[MPLS/0] 2d 02:19:12, metric 1
                      Receive
2                  *[MPLS/0] 2d 02:19:12, metric 1
                      Receive
800017             *[VPLS/7] 1d 13:54:49
                    > via vt-3/2/0.32769, Pop
800018             *[VPLS/7] 1d 13:54:59
                    > via vt-3/2/0.32772, Pop
vt-3/2/0.32769      [VPLS/7] 1d 13:54:49
                      Unusable
vt-3/2/0.32772      [VPLS/7] 1d 13:54:59
                      Unusable
```

 Copyright © 2018, Juniper Networks, Inc.

show route aspath-regex

List of Syntax

Syntax on page 2203
Syntax (EX Series Switches) on page 2203

Syntax

show route aspath-regex *regular-expression*
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route aspath-regex *regular-expression*

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display the entries in the routing table that match the specified autonomous system (AS) path regular expression.

Options

regular-expression—Regular expression that matches an entire AS path.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information

You can specify a regular expression as:

- An individual AS number
- A period wildcard used in place of an AS number
- An AS path regular expression that is enclosed in parentheses

You also can include the operators described in the table of AS path regular expression operators in the *Junos Policy Framework Configuration Guide*. The following list summarizes these operators:

- **{*m*,*n*}**—At least *m* and at most *n* repetitions of the AS path term.
- **{*m*}**—Exactly *m* repetitions of the AS path term.
- **{*m*,}**—*m* or more repetitions of the AS path term.
- *****—Zero or more repetitions of an AS path term.
- ***+***—One or more repetitions of an AS path term.
- ***?***—Zero or one repetition of an AS path term.
- ***aspath_term* | *aspath_term***—Match one of the two AS path terms.

  When you specify more than one AS number or path term, or when you include an operator in the regular expression, enclose the entire regular expression in quotation marks. For example, to match any path that contains AS number 234, specify the following command:

show route aspath-regex ".* 234 .*"

Required Privilege Level view

Related Documentation
- *Example: Using AS Path Regular Expressions*

List of Sample Output
show route aspath-regex (Matching a Specific AS Number) on page 2204
show route aspath-regex (Matching Any Path with Two AS Numbers) on page 2204

Output Fields For information about output fields, see the output field table for the show route command.

Sample Output

show route aspath-regex (Matching a Specific AS Number)

```
user@host> show route aspath-regex 65477
inet.0: 46411 destinations, 46411 routes (46409 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

111.222.1.0/25      *[BGP/170] 00:08:48, localpref 100, from 111.222.2.24
                    AS Path: [65477] ({65548 65536}) IGP
                      to 111.222.18.225 via fpa0.0(111.222.18.233)
111.222.1.128/25    *[IS-IS/15] 09:15:37, metric 37, tag 1
                      to 111.222.18.225 via fpa0.0(111.222.18.233)
                     [BGP/170] 00:08:48, localpref 100, from 111.222.2.24
                    AS Path: [65477] ({65548 65536}) IGP
                      to 111.222.18.225 via fpa0.0(111.222.18.233)
...
```

show route aspath-regex (Matching Any Path with Two AS Numbers)

```
user@host> show route aspath-regex ".* 234 3561 .*"

inet.0: 46351 destinations, 46351 routes (46349 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

9.20.0.0/17         *[BGP/170] 01:35:00, localpref 100, from 131.103.20.49
                    AS Path: [666] 234 3561 2685 2686 Incomplete
                      to 192.156.169.1 via 192.156.169.14(so-0/0/0)
12.10.231.0/24      *[BGP/170] 01:35:00, localpref 100, from 131.103.20.49
                    AS Path: [666] 234 3561 5696 7369 IGP
                      to 192.156.169.1 via 192.156.169.14(so-0/0/0)
24.64.32.0/19       *[BGP/170] 01:34:59, localpref 100, from 131.103.20.49
                    AS Path: [666] 234 3561 6327 IGP
                      to 192.156.169.1 via 192.156.169.14(so-0/0/0)
...
```

 Copyright © 2018, Juniper Networks, Inc.

show route best

List of Syntax

Syntax on page 2205
Syntax (EX Series Switches) on page 2205

Syntax

show route best *destination-prefix*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route best *destination-prefix*
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display the route in the routing table that is the best route to the specified address or range of addresses. The best route is the longest matching route.

Options

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

destination-prefix—Address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- show route brief on page 2208
- show route detail on page 2220

List of Sample Output

show route best on page 2205
show route best detail on page 2206
show route best extensive on page 2207
show route best terse on page 2207

Output Fields

For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route best

```
user@host> show route best 10.255.70.103
```

```
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32   *[OSPF/10] 1d 13:19:20, metric 2
                    > to 10.31.1.6 via ge-3/1/0.0
                      via so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32   *[RSVP/7] 1d 13:20:13, metric 2
                    > via so-0/3/0.0, label-switched-path green-r1-r3

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.0.0.0/8         *[Direct/0] 2d 01:43:34
                    > via fxp2.0
                    [Direct/0] 2d 01:43:34
                    > via fxp1.0
```

show route best detail

```
user@host> show route best 10.255.70.103 detail
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
        *OSPF   Preference: 10
                Next-hop reference count: 9
                Next hop: 10.31.1.6 via ge-3/1/0.0, selected
                Next hop: via so-0/3/0.0
                State: <Active Int>
                Local AS:    69
                Age: 1d 13:20:06        Metric: 2
                Area: 0.0.0.0
                Task: OSPF
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 5
                Next hop: via so-0/3/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 100016
                State: <Active Int>
                Local AS:    69
                Age: 1d 13:20:59        Metric: 2
                Task: RSVP
                Announcement bits (1): 1-Resolve tree 2
                AS path: I

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
10.0.0.0/8 (2 entries, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via fxp2.0, selected
                State: <Active Int>
```

 Copyright © 2018, Juniper Networks, Inc.

```
                Age: 2d 1:44:20
                Task: IF
                AS path: I
         Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via fxp1.0, selected
                State: <NotBest Int>
                Inactive reason: No difference
                Age: 2d 1:44:20
                Task: IF
                AS path: I
```

show route best extensive

The output for the **show route best extensive** command is identical to that for the **show route best detail** command. For sample output, see show route best detail on page 2206.

show route best terse

```
user@host> show route best 10.255.70.103 terse
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.255.70.103/32   O  10          2            >10.31.1.6
                                                  so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.255.70.103/32   R   7          2            >so-0/3/0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.0.0.0/8         D   0                       >fxp2.0
                     D   0                       >fxp1.0
```

Copyright © 2018, Juniper Networks, Inc.

show route brief

List of Syntax

Syntax on page 2208
Syntax (EX Series Switches) on page 2208

Syntax

show route brief
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route brief
<*destination-prefix*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display brief information about the active entries in the routing tables.

Options

none—Display all active entries in the routing table.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- show route all on page 2201
- show route best on page 2205

List of Sample Output

show route brief on page 2208

Output Fields

For information about output fields, see the Output Field table of the show route command.

Sample Output

show route brief

```
user@host> show route brief
inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 1w5d 20:30:29
                      Discard
10.255.245.51/32   *[Direct/0] 2w4d 13:11:14
                    > via lo0.0
```


Copyright © 2018, Juniper Networks, Inc.

```
172.16.0.0/12      *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
192.168.0.0/18      *[Static/5] 1w5d 20:30:29
                    > to 192.168.167.254 via fxp0.0
192.168.40.0/22    *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
192.168.64.0/18    *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
192.168.164.0/22   *[Direct/0] 2w4d 13:11:14
                    > via fxp0.0
192.168.164.51/32  *[Local/0] 2w4d 13:11:14
                      Local via fxp0.0
207.17.136.192/32  *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
green.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
100.101.0.0/16     *[Direct/0] 1w5d 20:30:28
                    > via fe-0/0/3.0
100.101.2.3/32     *[Local/0] 1w5d 20:30:28
                      Local via fe-0/0/3.0
172.16.233.5/32       *[OSPF/10] 1w5d 20:30:29, metric 1
                      MultiRecv
```

show route ccc

Syntax

show route ccc *ccc*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display circuit cross-connect (CCC) entries in the Multiprotocol Link Switching (MPLS) routing table.

Options

ccc—Name of an entry with a circuit cross-connect interface.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- *show connections*

List of Sample Output

show route ccc extensive on page 2210

Output Fields

For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route ccc extensive

```
user@host> show route ccc fe-0/1/0.600 extensive
mpls.0: 19 destinations, 19 routes (19 active, 0 holddown, 0 hidden)
fe-0/1/2.600 (1 entry, 1 announced)
TSI:
KRT in-kernel fe-0/1/2.600.0      /16 -> {0.0.0.0}
        *CCC    Preference: 7
                Next-hop reference count: 2
                Next hop: via so-0/0/3.0 weight 0x1, selected
                Label operation: Push 101424
                State: <Active Int>
                Local AS:   100
                Age: 28:13   Metric: 3
                Task: MPLS
                Announcement bits (1): 0-KRT
                AS path: I
```


Copyright © 2018, Juniper Networks, Inc.

show route community

List of Syntax

Syntax on page 2211
Syntax (EX Series Switches) on page 2211

Syntax

show route community *as-number:community-value*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route community *as-number:community-value*
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display the route entries in each routing table that are members of a Border Gateway Protocol (BGP) community.

Options

as-number:community-value—One or more community identifiers. ***as-number*** is the AS number, and ***community-value*** is the community identifier. When you specify more than one community identifier, enclose the identifiers in double quotation marks. Community identifiers can include wildcards.

For example:

user@host> **show route table inet.0 protocol bgp community "12083:6015" community "12083:65551"**

or

user@host> **show route table inet.0 protocol bgp community [12083:6014 12083:65551]**

brief | detail | extensive | terse—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information

Specifying the community option displays all routes matching the community found within the routing table. The community option does not limit the output to only the routes being advertised to the neighbor after any egress routing policy.

Required Privilege Level

view

Related Documentation

- show route detail on page 2220

List of Sample Output show route community on page 2212

Output Fields For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route community

```
user@host> show route community 234:80
inet.0: 46511 destinations, 46511 routes (46509 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.4.0/8          *[BGP/170] 03:33:07, localpref 100, from 131.103.20.49
                   AS Path: {666} 234 2548 1 IGP
                     to 192.156.169.1 via 192.156.169.14(so-0/0/0)
172.16.6.0/8          *[BGP/170] 03:33:07, localpref 100, from 131.103.20.49
                   AS Path: {666} 234 2548 568 721 Incomplete
                     to 192.156.169.1 via 192.156.169.14(so-0/0/0)
172.16.92.0/16         *[BGP/170] 03:33:06, localpref 100, from 131.103.20.49
                   AS Path: {666} 234 2548 1673 1675 1747 IGP
                     to 192.156.169.1 via 192.156.169.14(so-0/0/0)
```

 Copyright © 2018, Juniper Networks, Inc.

show route community-name

List of Syntax
Syntax on page 2213
Syntax (EX Series Switches) on page 2213

Syntax
show route community-name *community-name*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route community-name *community-name*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display the route entries in each routing table that are members of a Border Gateway Protocol (BGP) community, specified by a community name.

Options
community-name—Name of the community.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route community-name on page 2213

Output Fields
For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route community-name

```
user@host> show route community-name red-com
inet.0: 17 destinations, 17 routes (16 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

instance1.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.245.212/32   *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
                      AS path: 300 I
```

Copyright © 2018, Juniper Networks, Inc.

```
                      > to 172.16.100.1 via ge-1/1/0.0, label-switched-path to_fix
172.16.20.20/32       *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
                        AS path: I
                      > to 172.16.100.1 via ge-1/1/0.0, label-switched-path to_fix
172.16.100.0/24         *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
                        AS path: I
                      > to 172.16.100.1 via ge-1/1/0.0, label-switched-path to_fix

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.245.204:10:10.255.245.212/32
                     *[BGP/170] 00:06:40, localpref 100, from 10.255.245.204
                        AS path: 300 I
                      > to 172.16.100.1 via ge-1/1/0.0, label-switched-path to_fix
10.255.245.204:10:172.16.20.20/32
                     *[BGP/170] 00:36:02, localpref 100, from 10.255.245.204
                        AS path: I
                      > to 172.16.100.1 via ge-1/1/0.0, label-switched-path to_fix
10.255.245.204:10:100.1.4.0/24
                     *[BGP/170] 00:36:02, localpref 100, from 10.255.245.204
                        AS path: I
                      > to 172.16.100.1 via ge-1/1/0.0, label-switched-path to_fix

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

instance1.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

 Copyright © 2018, Juniper Networks, Inc.

show route damping

List of Syntax

Syntax on page 2215
Syntax (EX Series Switch and QFX Series) on page 2215

Syntax

show route damping (decayed | history | suppressed)
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and QFX Series)

show route damping (decayed | history | suppressed)
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display the BGP routes for which updates might have been reduced because of route flap damping.

Options

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.

decayed—Display route damping entries that might no longer be valid, but are not suppressed.

history—Display entries that have already been withdrawn, but have been logged.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

suppressed—Display entries that have been suppressed and are no longer being installed into the forwarding table or exported by routing protocols.

Required Privilege Level

view

Related Documentation

- clear bgp damping on page 1664
- show policy damping on page 1710

List of Sample Output

show route damping decayed detail on page 2218
show route damping history on page 2219
show route damping history detail on page 2219

Output Fields

Table 195 on page 2216 lists the output fields for the **show route damping** command. Output fields are listed in the approximate order in which they appear.

Table 195: show route damping Output Fields

Field Name	Field Description	Level of Output
routing-table-name	Name of the routing table—for example, **inet.0**.	All levels
destinations	Number of destinations for which there are routes in the routing table.	All levels
number **routes**	Number of routes in the routing table and total number of routes in the following states: • **active** • **holddown** (routes that are in a pending state before being declared inactive) • **hidden** (the routes are not used because of a routing policy)	All levels
destination-prefix **(entry, announced)**	Destination prefix. The **entry** value is the number of routes for this destination, and the **announced** value is the number of routes being announced for this destination.	**detail extensive**
[*protocol***,** *preference***]**	Protocol from which the route was learned and the preference value for the route. • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.	All levels
Next-hop reference count	Number of references made to the next hop.	**detail extensive**
Source	IP address of the route source.	**detail extensive**
Next hop	Network layer address of the directly reachable neighboring system.	**detail extensive**
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word **Selected**.	**detail extensive**
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.	**detail extensive**
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.	**detail extensive**
State	Flags for this route. For a description of possible values for this field, see the output field table for the show route detail command.	**detail extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 195: show route damping Output Fields (continued)

Field Name	Field Description	Level of Output
Local AS	AS number of the local routing device.	**detail extensive**
Peer AS	AS number of the peer routing device.	**detail extensive**
Age	How long the route has been known.	**detail extensive**
Metric	Metric for the route.	**detail extensive**
Task	Name of the protocol that has added the route.	**detail extensive**
Announcement bits	List of protocols that announce this route. ***n*-Resolve inet** indicates that the route is used for route resolution for next hops found in the routing table. ***n*** is an index used by Juniper Networks customer support only.	**detail extensive**
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device or if AS path prepending is configured. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.	All levels
to	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.	**brief** none
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word **Selected**.	**brief** none
Communities	Community path attribute for the route. See the output field table for the show route detail command.	**detail extensive**
Localpref	Local preference value included in the route.	All levels
Router ID	BGP router ID as advertised by the neighbor in the open message.	**detail extensive**

Table 195: show route damping Output Fields (continued)

Field Name	Field Description	Level of Output
Merit (last update/now)	Last updated and current figure-of-merit value.	**detail extensive**
damping-parameters	Name that identifies the damping parameters used, which is defined in the damping statement at the **[edit policy-options]** hierarchy level.	**detail extensive**
Last update	Time of most recent change in path attributes.	**detail extensive**
First update	Time of first change in path attributes, which started the route damping process.	**detail extensive**
Flaps	Number of times the route has gone up or down or its path attributes have changed.	**detail extensive**
Suppressed	(**suppressed** keyword only) This route is currently suppressed. A suppressed route does not appear in the forwarding table and routing protocols do not export it.	All levels
Reusable in	(**suppressed** keyword only) Time when a suppressed route will again be available.	All levels
Preference will be	(**suppressed** keyword only) Preference value that will be applied to the route when it is again active.	All levels

Sample Output

show route damping decayed detail

```
user@host> show route damping decayed detail
inet.0: 173319 destinations, 1533668 routes (172625 active, 4 holddown, 108083
hidden)
10.0.111.0/24 (7 entries, 1 announced)
        *BGP    Preference: 170/-101
                Next-hop reference count: 151973
                Source: 172.23.2.129
                Next hop: via so-1/2/0.0
                Next hop: via so-5/1/0.0, selected
                Next hop: via so-6/0/0.0
                Protocol next hop: 172.23.2.129
                Indirect next hop: 89a1a00 264185
                State: <Active Ext>
                Local AS: 64500 Peer AS:   64490
                Age: 3:28       Metric2: 0
                Task: BGP_64490.172.23.2.129+179
                Announcement bits (6): 0-KRT 1-RT 4-KRT 5-BGP.0.0.0.0+179

         6-Resolve tree 2 7-Resolve tree 3
                AS path: 64499 64510 645511 645511 645511 645511 I ()
                Communities: 65551:390 65551:2000 65551:3000 65550:701
                Localpref: 100
                Router ID: 172.23.2.129
                Merit (last update/now): 1934/1790
                damping-parameters: damping-high
```

 Copyright © 2018, Juniper Networks, Inc.

```
                Last update:       00:03:28 First update:       00:06:40
                Flaps: 2
```

show route damping history

```
user@host> show route damping history
inet.0: 173320 destinations, 1533529 routes (172624 active, 6 holddown, 108122
hidden)
+ = Active Route, - = Last Active, * = Both

10.108.0.0/15       [BGP ] 2d 22:47:58, localpref 100
                      AS path: 64220 65541 65542 I
                    > to 192.168.60.85 via so-3/1/0.0
```

show route damping history detail

```
user@host> show route damping history detail
inet.0: 173319 destinations, 1533435 routes (172627 active, 2 holddown, 108105
hidden)
10.108.0.0/15 (3 entries, 1 announced)
         BGP                 /-101
                Next-hop reference count: 69058
                Source: 192.168.60.85
                Next hop: 192.168.60.85 via so-3/1/0.0, selected
                State: <Hidden Ext>
                Inactive reason: Unusable path
                Local AS: 64500 Peer AS:  64220
                Age: 2d 22:48:10
                Task: BGP_64220.192.168.60.85+179
                AS path: 64220 65541 65542 I ()
                Communities: 65541:390 65541:2000 65541:3000 65504:3561
                Localpref: 100
                Router ID: 192.168.80.25
                Merit (last update/now): 1000/932
                damping-parameters: set-normal
                Last update:       00:01:05 First update:       00:01:05
                Flaps: 1
```

Copyright © 2018, Juniper Networks, Inc.

show route detail

List of Syntax

Syntax on page 2220
Syntax (EX Series Switches) on page 2220

Syntax

show route detail
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route detail
<*destination-prefix*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Comand introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display detailed information about the active entries in the routing tables.

Options

none—Display all active entries in the routing table on all systems.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show route detail on page 2231
show route detail (with BGP Multipath) on page 2237
show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 2238
show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 2238

Output Fields

Table 196 on page 2220 describes the output fields for the **show route detail** command. Output fields are listed in the approximate order in which they appear.

Table 196: show route detail Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).
number **destinations**	Number of destinations for which there are routes in the routing table.

Copyright © 2018, Juniper Networks, Inc.

Table 196: show route detail Output Fields (continued)

Field Name	Field Description
number **routes**	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active) • **holddown** (routes that are in the pending state before being declared inactive) • **hidden** (routes that are not used because of a routing policy)

route-destination **(entry, announced)**	Route destination (for example:10.0.0.1/24). The **entry** value is the number of routes for this destination, and the **announced** value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:
	• *MPLS-label* (for example, 80001).
	• *interface-name* (for example, ge-1/0/2).
	• *neighbor-address*:*control-word-status*:*encapsulation type*:*vc-id*:*source* (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96).
	  ◦ *neighbor-address*—Address of the neighbor.
	  ◦ *control-word-status*—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**.
	  ◦ *encapsulation type*—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.
	  ◦ *vc-id*—Virtual circuit identifier.
	  ◦ *source*—Source of the advertisement: **Local** or **Remote**.
	  ◦ *source*—Source of the advertisement: **Local** or **Remote**.

label stacking	(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels. • **S=0 route** indicates that a packet with an incoming label stack depth of 2 or more exits this routing device with one fewer label (the label-popping operation is performed). • If there is no **S=** information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).

Table 196: show route detail Output Fields (continued)

Field Name	Field Description

[*protocol*, *preference*]	Protocol from which the route was learned and the preference value for the route.
	• **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.
	• **-** —A hyphen indicates the last active route.
	• *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route.
	In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value.
	Preference2 values are signed integers, that is, **Preference2** values can be either positive or negative values. However, Junos OS evaluates **Preference2** values as unsigned integers that are represented by positive values. Based on the **Preference2** values, Junos OS evaluates a preferred route differently in the following scenarios:
	• **Both Signed Preference2 values**
	  ◦ Route A = -101
	  ◦ Route B = -156
	  Where both the **Preference2** values are signed, Junos OS evaluates only the unsigned value of **Preference2** and Route A, which has a lower **Preference2** value is preferred.
	• **Unsigned Preference2 values**
	  Now consider both unsigned **Preference2** values:
	  ◦ Route A = 4294967096
	  ◦ Route B = 200
	  Here, Junos OS considers the lesser **Preference2** value and Route B with a **Preference2** value of 200 is preferred because it is less than 4294967096.
	• **Combination of signed and unsigned Preference2 values**
	  When **Preference2** values of two routes are compared, and for one route the **Preference2** is a signed value, and for the other route it is an unsigned value, Junos OS prefers the route with the positive **Preference2** value over the negative **Preference2** value. For example, consider the following signed and unsigned **Preference2** values:
	  ◦ Route A = -200
	  ◦ Route B = 200
	  In this case, Route B with a **Preference2** value of 200 is preferred although this value is greater than -200, because Junos OS evaluates only the unsigned value of the **Preference2** value.

Level	(IS-IS only). In IS-IS, a single AS can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.
Route Distinguisher	IP subnet augmented with a 64-bit prefix.
PMSI	Provider multicast service interface (MVPN routing table).

Copyright © 2018, Juniper Networks, Inc.

Table 196: show route detail Output Fields (continued)

Field Name	Field Description
Next-hop type	Type of next hop. For a description of possible values for this field, see Table 197 on page 2226.
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word **Selected**. This field can also contain the following information: • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.
Label-switched-path ***lsp-path-name***	Name of the LSP used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label).
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.
State	State of the route (a route can be in more than one state). See Table 198 on page 2228.
Local AS	AS number of the local routing device.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.

Copyright © 2018, Juniper Networks, Inc.

Table 196: show route detail Output Fields (continued)

Field Name	Field Description
Metric*n*	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances. For sample output, see show route table.
Task	Name of the protocol that has added the route.
Announcement bits	The number of BGP peers or protocols to which Junos OS has announced this route, followed by the list of the recipients of the announcement. Junos OS can also announce the route to the KRT for installing the route into the Packet Forwarding Engine, to a resolve tree, a L2 VC, or even a VPN. For example, ***n*-Resolve inet** indicates that the specified route is used for route resolution for next hops found in the routing table. • ***n***—An index used by Juniper Networks customer support only.
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **Recorded**—The AS path is recorded by the sample process (sampled). • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used in the AS-path merge process, as defined in RFC 4893. • **[ ]**—If more than one AS number is configured on the routing device, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.

 Copyright © 2018, Juniper Networks, Inc.

Table 196: show route detail Output Fields (continued)

Field Name	Field Description
validation-state	(BGP-learned routes) Validation status of the route: • **Invalid**—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • **Unknown**—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • **Unverified**—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers. • **Valid**—Indicates that the prefix and autonomous system pair are found in the database.
ORR Generation-ID	Displays the optimal route reflection (ORR) generation identifier. ISIS and OSPF interior gateway protocol (IGP) updates filed whenever any of the corresponding ORR route has its metric valued changed, or if the ORR route is added or deleted.
FECs bound to route	Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
Primary Upstream	When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.
RPF Nexthops	When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.
Label	Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.
weight	Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Prefixes bound to route	Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See Table 199 on page 2230 for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: **none** or **Site Down**.
mtu	Maximum transmission unit (MTU) information.

Table 196: show route detail Output Fields (continued)

Field Name	Field Description
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).
Accepted Multipath	Current active path when BGP multipath is configured.
Accepted LongLivedStale	The LongLivedStale flag indicates that the route was marked LLGR-stale by this router, as part of the operation of LLGR receiver mode. Either this flag or the LongLivedStaleImport flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag.
Accepted LongLivedStaleImport	The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy. Either this flag or the LongLivedStale flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag. Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and import into the inet.0 routing table
ImportAccepted LongLivedStaleImport	Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and imported into the inet.0 routing table The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy.
Accepted MultipathContrib	Path currently contributing to BGP multipath.
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.

Table 197 on page 2226 describes all possible values for the Next-hop Types output field.

Table 197: Next-hop Types Output Field Values

Next-Hop Type	Description
Broadcast (bcast)	Broadcast next hop.
Deny	Deny next hop.
Discard	Discard next hop.
Dynamic List	Dynamic list next hop

 Copyright © 2018, Juniper Networks, Inc.

Table 197: Next-hop Types Output Field Values (continued)

Next-Hop Type	Description
Flood	Flood next hop. Consists of components called branches, up to a maximum of 32 branches. Each flood next-hop branch sends a copy of the traffic to the forwarding interface. Used by point-to-multipoint RSVP, point-to-multipoint LDP, point-to-multipoint CCC, and multicast.
Hold	Next hop is waiting to be resolved into a unicast or multicast type.
Indexed (idxd)	Indexed next hop.
Indirect (indr)	Used with applications that have a protocol next hop address that is remote. You are likely to see this next-hop type for internal BGP (IBGP) routes when the BGP next hop is a BGP neighbor that is not directly connected.
Interface	Used for a network address assigned to an interface. Unlike the router next hop, the interface next hop does not reference any specific node on the network.
Local (locl)	Local address on an interface. This next-hop type causes packets with this destination address to be received locally.
Multicast (mcst)	Wire multicast next hop (limited to the LAN).
Multicast discard (mdsc)	Multicast discard.
Multicast group (mgrp)	Multicast group member.
Receive (recv)	Receive.
Reject (rjct)	Discard. An ICMP unreachable message was sent.
Resolve (rslv)	Resolving next hop.
Routed multicast (mcrt)	Regular multicast next hop.
Router	A specific node or set of nodes to which the routing device forwards packets that match the route prefix. To qualify as next-hop type router, the route must meet the following criteria: • Must not be a direct or local subnet for the routing device. • Must have a next hop that is directly connected to the routing device.
Table	Routing table next hop.

Copyright © 2018, Juniper Networks, Inc.

Table 197: Next-hop Types Output Field Values (continued)

Next-Hop Type	Description
Unicast (ucst)	Unicast.
Unilist (ulst)	List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.

Table 198 on page 2228 describes all possible values for the State output field. A route can be in more than one state (for example, **<Active NoReadvrt Int Ext>**).

Table 198: State Output Field Values

Value	Description
Accounting	Route needs accounting.
Active	Route is active.
Always Compare MED	Path with a lower multiple exit discriminator (MED) is available.
AS path	Shorter AS path is available.
Cisco Non-deterministic MED selection	Cisco nondeterministic MED is enabled, and a path with a lower MED is available.
Clone	Route is a clone.
Cluster list length	Length of cluster list sent by the route reflector.
Delete	Route has been deleted.
Ex	Exterior route.
Ext	BGP route received from an external BGP neighbor.
FlashAll	Forces all protocols to be notified of a change to any route, active or inactive, for a prefix. When not set, protocols are informed of a prefix only when the active route changes.
Hidden	Route not used because of routing policy.
IfCheck	Route needs forwarding RPF check.
IGP metric	Path through next hop with lower IGP metric is available.
Inactive reason	Flags for this route, which was not selected as best for a particular destination.
Initial	Route being added.

 Copyright © 2018, Juniper Networks, Inc.

Table 198: State Output Field Values (continued)

Value	Description
Int	Interior route.
Int Ext	BGP route received from an internal BGP peer or a BGP confederation peer.
Interior > Exterior > Exterior via Interior	Direct, static, IGP, or EBGP path is available.
Local Preference	Path with a higher local preference value is available.
Martian	Route is a martian (ignored because it is obviously invalid).
MartianOK	Route exempt from martian filtering.
Next hop address	Path with lower metric next hop is available.
No difference	Path from neighbor with lower IP address is available.
NoReadvrt	Route not to be advertised.
NotBest	Route not chosen because it does not have the lowest MED.
Not Best in its group	Incoming BGP AS is not the best of a group (only one AS can be the best).
NotInstall	Route not to be installed in the forwarding table.
Number of gateways	Path with a greater number of next hops is available.
Origin	Path with a lower origin code is available.
Pending	Route pending because of a hold-down configured on another route.
Programmed	Route installed programatically by on-box or off-box applications using API.
Release	Route scheduled for release.
RIB preference	Route from a higher-numbered routing table is available.
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.
Route Metric or MED comparison	Route with a lower metric or MED is available.
Route Preference	Route with lower preference value is available
Router ID	Path through a neighbor with lower ID is available.

Table 198: State Output Field Values (continued)

Value	Description
Secondary	Route not a primary route.
Unusable path	Path is not usable because of one of the following conditions: • The route is damped. • The route is rejected by an import policy. • The route is unresolved.
Update source	Last tiebreaker is the lowest IP address value.
ProtectionCand	Indicates paths requesting protection.
ProtectionPath	Indicates the route entry that can be used as a protection path.

Table 199 on page 2230 describes the possible values for the Communities output field.

Table 199: Communities Output Field Values

Value	Description
area-number	4 bytes, encoding a 32-bit area number. For AS-external routes, the value is **0**. A nonzero value identifies the route as internal to the OSPF domain, and as within the identified area. Area numbers are relative to a particular OSPF domain.
bandwidth: ***local AS number:link-bandwidth-number***	Link-bandwidth community value used for unequal-cost load balancing. When BGP has several candidate paths available for multipath purposes, it does not perform unequal-cost load balancing according to the link-bandwidth community unless all candidate paths have this attribute.
domain-id	Unique configurable number that identifies the OSPF domain.
domain-id-vendor	Unique configurable number that further identifies the OSPF domain.
link-bandwidth-number	Link-bandwidth number: from **0** through **4,294,967,295** (bytes per second).
local AS number	Local AS number: from **1** through **65,535**.
options	1 byte. Currently this is only used if the route type is **5** or **7**. Setting the least significant bit in the field indicates that the route carries a type 2 metric.
origin	(Used with VPNs) Identifies where the route came from.
ospf-route-type	1 byte, encoded as **1** or **2** for intra-area routes (depending on whether the route came from a type 1 or a type 2 LSA); **3** for summary routes; **5** for external routes (area number must be**0**); **7** for NSSA routes; or **129** for sham link endpoint addresses.
route-type-vendor	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute **0x8000**. The format is ***area-number:ospf-route-type:options***.

 Copyright © 2018, Juniper Networks, Inc.

Table 199: Communities Output Field Values (continued)

Value	Description
rte-type	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute **0x0306**. The format is ***area-number:ospf-route-type:options***.
target	Defines which VPN the route participates in; **target** has the format ***32-bit IP address:16-bit number***. For example, 10.19.0.0:100.
unknown IANA	Incoming IANA codes with a value between **0x1** and **0x7fff**. This code of the BGP extended community attribute is accepted, but it is not recognized.
unknown OSPF vendor community	Incoming IANA codes with a value above **0x8000**. This code of the BGP extended community attribute is accepted, but it is not recognized.

Sample Output

show route detail

```
user@host> show route detail

inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 1:31:43
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

10.31.1.0/30 (2 entries, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 2
                Next hop: via so-0/3/0.0, selected
                State: <Active Int>
                Local AS:    69
                Age: 1:30:17
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I
         OSPF   Preference: 10
                Next-hop reference count: 1
                Next hop: via so-0/3/0.0, selected
                State: <Int>
                Inactive reason: Route Preference
                Local AS:    69
                Age: 1:30:17    Metric: 1
                ORR Generation-ID: 1
      Area: 0.0.0.0
                Task: OSPF
                AS path: I
```

Copyright © 2018, Juniper Networks, Inc.

```
10.31.1.1/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 7
                Interface: so-0/3/0.0
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:30:20
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I

...

10.31.2.0/30 (1 entry, 1 announced)
        *OSPF   Preference: 10
                Next-hop reference count: 9
                Next hop: via so-0/3/0.0
                Next hop: 10.31.1.6 via ge-3/1/0.0, selected
                State: <Active Int>
                Local AS:    69
                Age: 1:29:56    Metric: 2
                Area: 0.0.0.0
                ORR Generation-ID: 1
      Task: OSPF
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

172.16.233.2/32 (1 entry, 1 announced)
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:45
                Task: PIM Recv
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

172.16.233.22/32 (1 entry, 1 announced)
        *IGMP   Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:43
                Task: IGMP
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 100096
```

 Copyright © 2018, Juniper Networks, Inc.

```
                State: <Active Int>
                Local AS:    69
                Age: 1:25:49    Metric: 2
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

10.255.71.238/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: via so-0/3/0.0 weight 0x1, selected
                Label-switched-path green-r1-r2
                State: <Active Int>
                Local AS:    69
                Age: 1:25:49    Metric: 1
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

private__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    69
                Age: 1:31:44
                Task: IF
                AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
0   (1 entry, 1 announced)
        *MPLS   Preference: 0
                Next hop type: Receive
                Next-hop reference count: 6
                State: <Active Int>
                Local AS:    69
                Age: 1:31:45    Metric: 1
                Task: MPLS
                Announcement bits (1): 0-KRT
                AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

  299840 (1 entry, 1 announced)
  TSI:
  KRT in-kernel 299840 /52 -> {indirect(1048575)}
          *RSVP   Preference: 7/2
                  Next hop type: Flood
                  Address: 0x9174a30
                  Next-hop reference count: 4
                  Next hop type: Router, Next hop index: 798
                  Address: 0x9174c28
                  Next-hop reference count: 2
```

```
                Next hop: 172.16.0.2 via lt-1/2/0.9 weight 0x1
                Label-switched-path R2-to-R4-2p2mp
                Label operation: Pop
                Next hop type: Router, Next hop index: 1048574
                Address: 0x92544f0
                Next-hop reference count: 2
                Next hop: 172.16.0.2 via lt-1/2/0.7 weight 0x1
                Label-switched-path R2-to-R200-p2mp
                Label operation: Pop
                Next hop: 172.16.0.2 via lt-1/2/0.5 weight 0x8001
                Label operation: Pop
                State: <Active Int>
                Age: 1:29       Metric: 1
                Task: RSVP
                Announcement bits (1): 0-KRT
                AS path: I...

800010 (1 entry, 1 announced)
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: via vt-3/2/0.32769, selected
                Label operation: Pop
                State: <Active Int>
                Age: 1:29:30
                Task: Common L2 VC
                Announcement bits (1): 0-KRT
                AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 800012, Push 100096(top)
                Protocol next hop: 10.255.70.103
                Push 800012
                Indirect next hop: 87272e4 1048574
                State: <Active Int>
                Age: 1:29:30    Metric2: 2
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 1-Common L2 VC
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

abcd::10:255:71:52/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    69
                Age: 1:31:44
                Task: IF
                AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
```

 Copyright © 2018, Juniper Networks, Inc.

```
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:44
                Task: IF
                AS path: I

ff02::2/128 (1 entry, 1 announced)
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:45
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::d/128 (1 entry, 1 announced)
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:45
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::16/128 (1 entry, 1 announced)
        *MLD    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:43
                Task: MLD
                Announcement bits (1): 0-KRT
                AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.16385, selected
                State: <Active NoReadvrt Int>
                Age: 1:31:44
                Task: IF
                AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

10.255.70.103:1:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.70.103:1
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    69 Peer AS:    69
```

Copyright © 2018, Juniper Networks, Inc.

```
                    Age: 1:25:49    Metric2: 1
                    AIGP 210
                    Task: BGP_69.10.255.70.103+179
                    Announcement bits (1): 0-green-l2vpn
                    AS path: I
                    Communities: target:11111:1 Layer2-info: encaps:VPLS,
                    control flags:, mtu: 0
                    Label-base: 800008, range: 8
                    Localpref: 100
                    Router ID: 10.255.70.103
                    Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
        *L2VPN  Preference: 170/-1
                Next-hop reference count: 5
                Protocol next hop: 10.255.71.52
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 1:31:40    Metric2: 1
                Task: green-l2vpn
                Announcement bits (1): 1-BGP.0.0.0.0+179
                AS path: I
                Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
                mtu: 0
                Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
        *L2VPN  Preference: 170/-101
                Next-hop reference count: 5
                Protocol next hop: 10.255.71.52
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 1:31:40    Metric2: 1
                Task: green-l2vpn
                Announcement bits (1): 1-BGP.0.0.0.0+179
                AS path: I
                Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
                Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                Next hop: via so-1/1/2.0 weight 1, selected
                Label-switched-path my-lsp
                Label operation: Push 100000[0]
                Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
                State: <Active Int>
                Local AS:    99
                Age: 10:21
                Task: l2 circuit
                Announcement bits (1): 0-LDP
                AS path: I
                VC Label 100000, MTU 1500, VLAN ID 512

inet.0: 45 destinations, 47 routes (44 active, 0 holddown, 1 hidden)
1.1.1.3/32 (1 entry, 1 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Router, Next hop index: 580
```


Copyright © 2018, Juniper Networks, Inc.

```
                Address: 0x9db6ed0
                Next-hop reference count: 8
                Next hop: 10.1.1.6 via lt-1/0/10.5, selected
                Session Id: 0x18a
                State: <Active Int>
                Local AS:     2
                Age: 1:32       Metric: 10
                Validation State: unverified
                ORR Generation-ID: 1
                Task: IS-IS
                Announcement bits (3): 0-KRT 5-Resolve tree 4 6-Resolve_IGP_FRR
task
                AS path: I

inet.0: 61 destinations, 77 routes (61 active, 1 holddown, 0 hidden)
1.1.1.1/32 (2 entries, 1 announced)
        *OSPF   Preference: 10
                Next hop type: Router, Next hop index: 673
                Address: 0xc008830
                Next-hop reference count: 3
                Next hop: 10.1.1.1 via ge-0/0/2.0, selected
                Session Id: 0x1b7
                State: <Active Int>
                Local AS:     1
                Age: 3:06:59    Metric: 100
                Validation State: unverified
                ORR Generation-ID: 1
                Area: 0.0.0.0
                Task: OSPF
                Announcement bits (2): 1-KRT 9-Resolve tree 2
                AS path: I
```

show route detail (with BGP Multipath)

```
user@host> show route detail

10.1.1.8/30 (2 entries, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 262142
                Address: 0x901a010
                Next-hop reference count: 2
                Source: 10.1.1.2
                Next hop: 10.1.1.2 via ge-0/3/0.1, selected
                Next hop: 10.1.1.6 via ge-0/3/0.5
                State: <Active Ext>
                Local AS:     1 Peer AS:     2
                Age: 5:04:43
               Validation State: unverified
                Task: BGP_2.10.1.1.2+59955
                Announcement bits (1): 0-KRT
                AS path: 2 I
                Accepted Multipath
                Localpref: 100
                Router ID: 172.16.1.2
         BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 678
                Address: 0x8f97520
                Next-hop reference count: 9
                Source: 10.1.1.6
                Next hop: 10.1.1.6 via ge-0/3/0.5, selected
                State: <NotBest Ext>
```

```
               Inactive reason: Not Best in its group - Active preferred
               Local AS:     1 Peer AS:     2
               Age: 5:04:43
            Validation State: unverified
               Task: BGP_2.10.1.1.6+58198
               AS path: 2 I
               Accepted MultipathContrib
               Localpref: 100
               Router ID: 172.16.1.3
```

show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Next-hop reference count: 3
                Address: 0x9097d90
                Next hop: via vt-0/1/0.1
                Next-hop index: 661
                Label operation: Pop
                Address: 0x9172130
                Next hop: via so-0/0/3.0
                Next-hop index: 654
                Label operation: Swap 299872
                State: **Active Int>
                Local AS:  1001
                Age: 8:20       Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
               FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
 src 192.168.142.2
```

show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show route label 301568 detail

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Address: 0x2735208
                Next-hop reference count: 3
                Next hop type: Router, Next hop index: 1397
                Address: 0x2735d2c
                Next-hop reference count: 3
                Next hop: 1.3.8.2 via ge-1/2/22.0
                Label operation: Pop
                Load balance label:  None;
                Next hop type: Router, Next hop index: 1395
                Address: 0x2736290
                Next-hop reference count: 3
                Next hop: 1.3.4.2 via ge-1/2/18.0
                Label operation: Pop
                Load balance label:  None;
                State: <Active Int AckRequest MulticastRPF>
                Local AS:    10
                Age: 54:05      Metric: 1
```

 Copyright © 2018, Juniper Networks, Inc.

```
                Validation State: unverified
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                FECs bound to route: P2MP root-addr 172.16.1.1, grp: 232.1.1.1,
src: 192.168.219.11
                Primary Upstream : 172.16.1.3:0--172.16.1.2:0
                   RPF Nexthops  :
                        ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
                        ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
                Backup Upstream : 172.16.1.3:0--172.16.1.6:0
                   RPF Nexthops :
                        ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xfffe
                        ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xfffe
```

Copyright © 2018, Juniper Networks, Inc.

show route exact

List of Syntax

Syntax on page 2240
Syntax (EX Series Switches) on page 2240

Syntax

show route exact *destination-prefix*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route exact *destination-prefix*
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display only the routes that exactly match the specified address or range of addresses.

Options

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

destination-prefix—Address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show route exact on page 2240
show route exact detail on page 2241
show route exact extensive on page 2241
show route exact terse on page 2241

Output Fields

For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route exact

```
user@host> show route exact 207.17.136.0/24

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
207.17.136.0/24    *[Static/5] 2d 03:30:22
                    > to 192.168.71.254 via fxp0.0
```

 Copyright © 2018, Juniper Networks, Inc.

show route exact detail

```
user@host> show route exact 207.17.136.0/24 detail

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
207.17.136.0/24 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2d 3:30:26
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I
```

show route exact extensive

```
user@host> show route exact 207.17.136.0/24 extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 1:25:18
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I
```

show route exact terse

```
user@host> show route exact 207.17.136.0/24 terse

inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 207.17.136.0/24    S   5                       >192.168.71.254
```

show route export

List of Syntax

Syntax on page 2242
Syntax (EX Series Switches) on page 2242

Syntax

show route export
<brief | detail>
<instance <*instance-name*> | *routing-table-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route export
<brief | detail>
<instance <*instance-name*> | *routing-table-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display policy-based route export information. Policy-based export simplifies the process of exchanging route information between routing instances.

Options

none—(Same as **brief**.) Display standard information about policy-based export for all instances and routing tables on all systems.

brief | detail—(Optional) Display the specified level of output.

instance <*instance-name*>—(Optional) Display a particular routing instance for which policy-based export is currently enabled.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

routing-table-name—(Optional) Display information about policy-based export for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the **show route export inet** command).

Required Privilege Level

view

List of Sample Output

show route export on page 2243
show route export detail on page 2243
show route export instance detail on page 2244

Output Fields

Table 200 on page 2243 lists the output fields for the **show route export** command. Output fields are listed in the approximate order in which they appear.

 Copyright © 2018, Juniper Networks, Inc.

Table 200: show route export Output Fields

Field Name	Field Description	Level of Output
Table or *table-name*	Name of the routing tables that either import or export routes.	All levels
Routes	Number of routes exported from this table into other tables. If a particular route is exported to different tables, the counter will only increment by one.	**brief** none
Export	Whether the table is currently exporting routes to other tables: **Y** or **N** (Yes or No).	**brief** none
Import	Tables currently importing routes from the originator table. (Not displayed for tables that are not exporting any routes.)	**detail**
Flags	(**instance** keyword only) Flags for this feature on this instance: • **config auto-policy**—The policy was deduced from the configured IGP export policies. • **cleanup**—Configuration information for this instance is no longer valid. • **config**—The instance was explicitly configured.	**detail**
Options	(**instance** keyword only) Configured option displays the type of routing tables the feature handles: • **unicast**—Indicates *instance*.inet.0. • **multicast**—Indicates *instance*.inet.2. • **unicast multicast**—Indicates *instance*.inet.0 and *instance*.inet.2.	**detail**
Import policy	(**instance** keyword only) Policy that **route export** uses to construct the import-export matrix. Not displayed if the instance type is **vrf**.	**detail**
Instance	(**instance** keyword only) Name of the routing instance.	**detail**
Type	(**instance** keyword only) Type of routing instance: **forwarding**, **non-forwarding**, or **vrf**.	**detail**

Sample Output

show route export

```
user@host> show route export
Table                             Export          Routes
inet.0                            N                    0
black.inet.0                      Y                    3
red.inet.0                        Y                    4
```

show route export detail

```
user@host>  show route export detail
inet.0                                          Routes:        0
black.inet.0                                    Routes:        3
  Import: [ inet.0 ]
red.inet.0                                      Routes:        4
  Import: [ inet.0 ]
```

show route export instance detail

```
user@host> show route export instance detail
Instance: master                        Type: forwarding
  Flags: <config auto-policy> Options: <unicast multicast>
  Import policy: [ (ospf-master-from-red || isis-master-from-black) ]
Instance: black                         Type: non-forwarding
Instance: red                           Type: non-forwarding
```


Copyright © 2018, Juniper Networks, Inc.

show route export vrf-target

Syntax

show route export vrf-target
<brief | detail>
<community *community--regular-expression*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display the VPN routing and forwarding (VRF) target communities for which policy-based route export is currently distributing routes. This command is relevant when there are overlapping virtual private networks (VPNs).

Options

none—Display standard information about all target communities.

brief | detail—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.

community *community-regular-expression*—(Optional) Display information about the specified community.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show route export vrf-target on page 2246
show route export vrf-target community on page 2246
show route export vrf-target detail on page 2246

Output Fields

Table 201 on page 2245 lists the output fields for the **show route export vrf-target** command. Output fields are listed in the approximate order in which they appear.

Table 201: show route export vrf-target Output Fields

Field Name	Field Description	Level of Output
Route target	Target communities for which auto-export is currently distributing routes.	**brief** none
Family	Routing table entries for the specified family.	**brief** none
type-of-routing-table(s)	Type of routing tables the feature handles: • **unicast**—Indicates ***instance*.inet.0**. • **multicast**—Indicates ***instance*.inet.2**. • **unicast multicast**—Indicates ***instance*.inet.0** and ***instance*.inet.2**.	**brief** none
Import	Number of routing tables that are currently importing routes with this target community. Omitted for tables that are not importing routes.	**brief** none

Copyright © 2018, Juniper Networks, Inc.

Table 201: show route export vrf-target Output Fields (continued)

Field Name	Field Description	Level of Output
Export	Number of routing tables that are currently exporting routes with this target community. Omitted for tables that are not exporting routes.	**brief** none
Target	Target communities, family, and options for which auto-export is currently distributing routes.	**detail**
Import table(s)	Name of the routing tables that are importing a particular route target.	**detail**
Export table(s)	Name of the routing tables that are exporting a particular route target.	**detail**

Sample Output

show route export vrf-target

```
user@host> show route export vrf-target
Route Target                 Family                  Import     Export
69:1                         inet      unicast            2          2
69:2                         inet      unicast            2          2
```

show route export vrf-target community

```
user@host> show route export vrf-target community target:69:1
Route Target                 Family                  Import     Export
69:1                         inet      unicast            2          2
```

show route export vrf-target detail

```
user@host> show route export vrf-target detail
Target: 1:12                              inet      unicast
  Import table(s): vrf-11.inet.0 vrf-12.inet.0
  Export table(s): vrf-12.inet.0
Target: 1:13                              inet      unicast
  Import table(s): vrf-12.inet.0 vrf-13.inet.0
  Export table(s): vrf-13.inet.0
```


Copyright © 2018, Juniper Networks, Inc.

show route extensive

List of Syntax

Syntax on page 2247
Syntax (EX Series Switches) on page 2247

Syntax

show route extensive
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route extensive
<*destination-prefix*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display extensive information about the active entries in the routing tables.

Options

none—Display all active entries in the routing table.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show route extensive on page 2254
show route extensive (Access Route) on page 2261
show route extensive (BGP PIC Edge) on page 2261
show route extensive (FRR and LFA) on page 2262
show route extensive (IS-IS) on page 2263
show route extensive (Route Reflector) on page 2263
show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 2263
show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 2264

Output Fields

Table 202 on page 2247 describes the output fields for the **show route extensive** command. Output fields are listed in the approximate order in which they appear.

Table 202: show route extensive Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).
number **destinations**	Number of destinations for which there are routes in the routing table.

Copyright © 2018, Juniper Networks, Inc.

Table 202: show route extensive Output Fields (continued)

Field Name	Field Description
number routes	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active). • **holddown** (routes that are in the pending state before being declared inactive). • **hidden** (routes that are not used because of a routing policy).

route-destination (entry, announced)	Route destination (for example: 10.0.0.1/24). The **entry** value is the number of route for this destination, and the **announced** value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:
	• ***MPLS-label*** (for example, 80001).
	• ***interface-name*** (for example, ge-1/0/2).
	• ***neighbor-address*:*control-word-status*:*encapsulation type*:*vc-id*:*source*** (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96).
	  ◦ ***neighbor-address***—Address of the neighbor.
	  ◦ ***control-word-status***—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**.
	  ◦ ***encapsulation type***—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.
	  ◦ ***vc-id***—Virtual circuit identifier.
	  ◦ ***source***—Source of the advertisement: **Local** or **Remote**.

TSI	Protocol header information.
label stacking	(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels. • **S=0 route** indicates that a packet with an incoming label stack depth of two or more exits this router with one fewer label (the label-popping operation is performed). • If there is no **S=** information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).
[*protocol*, *preference*]	Protocol from which the route was learned and the preference value for the route. • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.

 Copyright © 2018, Juniper Networks, Inc.

Table 202: show route extensive Output Fields (continued)

Field Name	Field Description
Level	(IS-IS only). In IS-IS, a single autonomous system (AS) can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.
Route Distinguisher	IP subnet augmented with a 64-bit prefix.
PMSI	Provider multicast service interface (MVPN routing table).
Next-hop type	Type of next hop. For a description of possible values for this field, see the Output Field table in the show route detail command.
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word **Selected**. This field can also contain the following information: • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.
Label-switched-path ***lsp-path-name***	Name of the LSP used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label).
Offset	Whether the metric has been increased or decreased by an offset value.
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to recursively derive a forwarding next hop.

Copyright © 2018, Juniper Networks, Inc.

Table 202: show route extensive Output Fields (continued)

Field Name	Field Description
label-operation	MPLS label and operation occurring at this routing device. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label).
Indirect next hops	When present, a list of nodes that are used to resolve the path to the next-hop destination, in the order that they are resolved. When BGP PIC Edge is enabled, the output lines that contain **Indirect next hop: weight** follow next hops that the software can use to repair paths where a link failure occurs. The next-hop weight has one of the following values: • 0x1 indicates active next hops. • 0x4000 indicates passive next hops.
State	State of the route (a route can be in more than one state). See the Output Field table in the show route detail command.
Session ID	The BFD session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA).
Weight	Weight for the backup path. If the weight of an indirect next hop is larger than zero, the weight value is shown. For sample output, see show route table.

 Copyright © 2018, Juniper Networks, Inc.

Table 202: show route extensive Output Fields (continued)

Field Name	Field Description
Inactive reason	If the route is inactive, the reason for its current state is indicated. Typical reasons include: • **Active preferred**—Currently active route was selected over this route. • **Always compare MED**—Path with a lower multiple exit discriminator (MED) is available. • **AS path**—Shorter AS path is available. • **Cisco Non-deterministic MED selection**—Cisco nondeterministic MED is enabled and a path with a lower MED is available. • **Cluster list length**—Path with a shorter cluster list length is available. • **Forwarding use only**—Path is only available for forwarding purposes. • **IGP metric**—Path through the next hop with a lower IGP metric is available. • **IGP metric type**—Path with a lower OSPF link-state advertisement type is available. • **Interior > Exterior > Exterior via Interior**—Direct, static, IGP, or EBGP path is available. • **Local preference**—Path with a higher local preference value is available. • **Next hop address**—Path with a lower metric next hop is available. • **No difference**—Path from a neighbor with a lower IP address is available. • **Not Best in its group**—Occurs when multiple peers of the same external AS advertise the same prefix and are grouped together in the selection process. When this reason is displayed, an additional reason is provided (typically one of the other reasons listed). • **Number of gateways**—Path with a higher number of next hops is available. • **Origin**—Path with a lower origin code is available. • **OSPF version**—Path does not support the indicated OSPF version. • **RIB preference**—Route from a higher-numbered routing table is available. • **Route destinguisher**—64-bit prefix added to IP subnets to make them unique. • **Route metric or MED comparison**—Route with a lower metric or MED is available. • **Route preference**—Route with a lower preference value is available. • **Router ID**—Path through a neighbor with a lower ID is available. • **Unusable path**—Path is not usable because of one of the following conditions: the route is damped, the route is rejected by an import policy, or the route is unresolved. • **Update source**—Last tiebreaker is the lowest IP address value.
Local AS	Autonomous system (AS) number of the local routing device.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.
Metric	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances. For sample output, see show route table.

Table 202: show route extensive Output Fields (continued)

Field Name	Field Description
Task	Name of the protocol that has added the route.
Announcement bits	List of protocols that are consumers of the route. Using the following output as an example, **Announcement bits (3): 0-KRT 5-Resolve tree 2 8-BGP RT Background** there are (3) announcement bits to reflect the three clients (protocols) that have state for this route: Kernel (0-KRT), 5 (resolution tree process 2), and 8 (BGP). The notation ***n*-Resolve inet** indicates that the route is used for route resolution for next hops found in the routing table. ***n*** is an index used by Juniper Networks customer support only.
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **Recorded**—The AS path is recorded by the sample process (sampled). • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.
validation-state	(BGP-learned routes) Validation status of the route: • **Invalid**—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • **Unknown**—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • **Unverified**—Indicates that origin validation is not enabled for the BGP peers. • **Valid**—Indicates that the prefix and autonomous system pair are found in the database.
FECs bound to route	Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
AS path: I <Originator>	(For route reflected output only) Originator ID attribute set by the route reflector.

 Copyright © 2018, Juniper Networks, Inc.

Table 202: show route extensive Output Fields (continued)

Field Name	Field Description
route status	Indicates the status of a BGP route: • **Accepted**—The specified BGP route is imported by the default BGP policy. • **Import**—The route is imported into a Layer 3 VPN routing instance. • **Import-Protect**—A remote instance egress that is protected. • **Multipath**—A BGP multipath active route. • **MultipathContrib**—The route is not active but contributes to the BGP multipath. • **Protect**—An egress route that is protected. • **Stale**—A route that is marked stale due to graceful restart.
Primary Upstream	When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.
RPF Nexthops	When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.
Label	Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.
weight	Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.
Originator ID	(For route reflected output only) Address of router that originally sent the route to the route reflector.
Prefixes bound to route	Forwarding Equivalent Class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See the Output Field table in the show route detail command for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: **none** or Site Down.
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.

Copyright © 2018, Juniper Networks, Inc.

Table 202: show route extensive Output Fields (continued)

Field Name	Field Description
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.
Originating RIB	Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of inet.0 resolving through inet.0 and inet.3, this field indicates which routing table, inet.0 or inet.3, provided the best path for a particular prefix.
Node path count	Number of nodes in the path.
Forwarding nexthops	Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.

Sample Output

show route extensive

```
user@host> show route extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
203.0.113.10/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 203.0.113.10/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    64496
                Age: 1:34:06
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

203.0.113.30/30 (2 entries, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 2
                Next hop: via so-0/3/0.0, selected
                State: <Active Int>
                Local AS:    64496
                Age: 1:32:40
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I
         OSPF   Preference: 10
                Next-hop reference count: 1
                Next hop: via so-0/3/0.0, selected
```


Copyright © 2018, Juniper Networks, Inc.

```
                State: <Int>
                Inactive reason: Route Preference
                Local AS:    64496
                Age: 1:32:40    Metric: 1
                Area: 0.0.0.0
                Task: OSPF
                AS path: I

203.0.113.103/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 7
                Interface: so-0/3/0.0
                State: <Active NoReadvrt Int>
                Local AS:    644969
                Age: 1:32:43
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I

...

203.0.113.203/30 (1 entry, 1 announced)
TSI:
KRT in-kernel 203.0.113.203/30 -> {203.0.113.216}
        *OSPF   Preference: 10
                Next-hop reference count: 9
                Next hop: via so-0/3/0.0
                Next hop: 203.0.113.216 via ge-3/1/0.0, selected
                State: <Active Int>
                Local AS:    64496
                Age: 1:32:19    Metric: 2
                Area: 0.0.0.0
                Task: OSPF
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

198.51.100.2/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 198.51.100.2/32 -> {}
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:08
                Task: PIM Recv
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

198.51.100.22/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 198.51.100.22/32 -> {}
        *IGMP   Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:06
```

```
                Task: IGMP
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

203.0.113.103/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: 203.0.113.216 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 100096
                State: <Active Int>
                Local AS:    64496
                Age: 1:28:12    Metric: 2
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

203.0.113.238/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: via so-0/3/0.0 weight 0x1, selected
                Label-switched-path green-r1-r2
                State: <Active Int>
                Local AS:    64496
                Age: 1:28:12    Metric: 1
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

...

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    64496
                Age: 1:34:07
                Task: IF
                AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

0   (1 entry, 1 announced)
TSI:
KRT in-kernel 0      /36 -> {}
        *MPLS   Preference: 0
                Next hop type: Receive
                Next-hop reference count: 6
                State: <Active Int>
                Local AS:    64496
                Age: 1:34:08    Metric: 1
```


Copyright © 2018, Juniper Networks, Inc.

```
                Task: MPLS
                Announcement bits (1): 0-KRT
                AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
  299840 (1 entry, 1 announced)
  TSI:
  KRT in-kernel 299840 /52 -> {indirect(1048575)}
          *RSVP   Preference: 7/2
                  Next hop type: Flood
                  Address: 0x9174a30
                  Next-hop reference count: 4
                  Next hop type: Router, Next hop index: 798
                  Address: 0x9174c28
                  Next-hop reference count: 2
                  Next hop: 198.51.100.2 via lt-1/2/0.9 weight 0x1
                  Label-switched-path R2-to-R4-2p2mp
                  Label operation: Pop
                  Next hop type: Router, Next hop index: 1048574
                  Address: 0x92544f0
                  Next-hop reference count: 2
                  Next hop: 198.51.100.2 via lt-1/2/0.7 weight 0x1
                  Label-switched-path R2-to-R200-p2mp
                  Label operation: Pop
                  Next hop: 198.51.100.2 via lt-1/2/0.5 weight 0x8001
                  Label operation: Pop
                  State: <Active Int>
                  Age: 1:29       Metric: 1
                  Task: RSVP
                  Announcement bits (1): 0-KRT
                  AS path: I...

800010 (1 entry, 1 announced)

TSI:
KRT in-kernel 800010 /36 -> {vt-3/2/0.32769}
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: via vt-3/2/0.32769, selected
                Label operation: Pop
                State: <Active Int>
                Age: 1:31:53
                Task: Common L2 VC
                Announcement bits (1): 0-KRT
                AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
TSI:
KRT in-kernel vt-3/2/0.32769.0      /16 -> {indirect(1048574)}
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: 203.0.113.216 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 800012, Push 100096(top)
                Protocol next hop: 203.0.113.103
                Push 800012
                Indirect next hop: 87272e4 1048574
                State: <Active Int>
                Age: 1:31:53    Metric2: 2
```

```
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 1-Common L2 VC
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Indirect next hops: 1
                        Protocol next hop: 203.0.113.103 Metric: 2
                        Push 800012
                        Indirect next hop: 87272e4 1048574
                        Indirect path forwarding next hops: 1
                               Next hop: 203.0.113.216 via ge-3/1/0.0 weight 0x1

                        203.0.113.103/32 Originating RIB: inet.3
                          Metric: 2                      Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: 203.0.113.216 via ge-3/1/0.0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

2001:db8::10:255:71:52/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    64496
                Age: 1:34:07
                Task: IF
                AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:07
                Task: IF
                AS path: I

ff02::2/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::2/128 -> {}
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:08
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::d/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::d/128 -> {}
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:08
```

 Copyright © 2018, Juniper Networks, Inc.

```
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::16/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::16/128 -> {}
        *MLD    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:06
                Task: MLD
                Announcement bits (1): 0-KRT
                AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.16385, selected
                State: <Active NoReadvrt Int>
                Age: 1:34:07
                Task: IF
                AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

203.0.113.103:1:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 203.0.113.103:1
                Next-hop reference count: 7
                Source: 203.0.113.103
                Protocol next hop: 203.0.113.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    64496 Peer AS:    64496
                Age: 1:28:12    Metric2: 1
                Task: BGP_69.203.0.113.103+179
                Announcement bits (1): 0-green-l2vpn
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Label-base: 800008, range: 8
                Localpref: 100
                Router ID: 203.0.113.103
                Primary Routing Table bgp.l2vpn.0

203.0.113.152:1:1:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699540
        *L2VPN  Preference: 170/-1
                Next-hop reference count: 5
                Protocol next hop: 203.0.113.152
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 1:34:03    Metric2: 1
                Task: green-l2vpn
                Announcement bits (1): 1-BGP.0.0.0.0+179
```

```
                AS path: I
                Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
                mtu: 0
                Label-base: 800016, range: 8, status-vector: 0x9F

203.0.113.152:1:5:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699528
        *L2VPN  Preference: 170/-101
                Next-hop reference count: 5
                Protocol next hop: 203.0.113.152
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 1:34:03    Metric2: 1
                Task: green-l2vpn
                Announcement bits (1): 1-BGP.0.0.0.0+179
                AS path: I
                Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
                Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

TSI:

203.0.113.163:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                Next hop: via so-1/1/2.0 weight 1, selected
                Label-switched-path my-lsp
                Label operation: Push 100000[0]
                Protocol next hop: 203.0.113.163 Indirect next hop: 86af000 296
                State: <Active Int>
                Local AS:    64499
                Age: 10:21
                Task: l2 circuit
                Announcement bits (1): 0-LDP
                AS path: I
                VC Label 100000, MTU 1500, VLAN ID 512


203.0.113.55/24 (1 entry, 1 announced)
TSI:
KRT queued (pending) add
  198.51.100.0/24 -> {Push 300112}
        *BGP    Preference: 170/-101
                Next hop type: Router
                Address: 0x925c208
                Next-hop reference count: 2
                Source: 203.0.113.9
                Next hop: 203.0.113.9 via ge-1/2/0.15, selected
                Label operation: Push 300112
                Label TTL action: prop-ttl
                State: <Active Ext>
                Local AS:  64509 Peer AS: 65539
                Age: 1w0d 23:06:56
                AIGP: 25
                Task: BGP_65539.203.0.113.9+56732
                Announcement bits (1): 0-KRT
                AS path: 65539 64508 I
                Accepted
```


Copyright © 2018, Juniper Networks, Inc.

```
                Route Label: 300112
                Localpref: 100
                Router ID: 213.0.113.99
```

show route extensive (Access Route)

```
user@host> show route 203.0.113.102 extensive
inet.0: 39256 destinations, 39258 routes (39255 active, 0 holddown, 1 hidden)
203.0.113.102/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 203.0.113.102/32 -> {192.0.2.2}
OSPF area : 0.0.0.0, LSA ID : 203.0.113.102, LSA type : Extern
        *Access Preference: 13
                Next-hop reference count: 78472
                Next hop: 192.0.2.2 via fe-0/0/0.0, selected
                State: <Active Int>
       Age: 12
                Task: RPD Unix Domain Server./var/run/rpd_serv.local
                Announcement bits (2): 0-KRT 1-OSPFv2
                AS path: I


user@host> show route 2001:db8:4641:1::/48 extensive

inet6.0: 75 destinations, 81 routes (75 active, 0 holddown, 0 hidden)
  2001:db8:4641:1::/48 (1 entry, 1 announced)
  TSI:
  KRT in-kernel 2001:db8:4641:1::/48 -> {#0 0.13.1.0.0.1}
          *Access Preference: 13
                  Next hop type: Router, Next hop index: 74548
                  Address: 0x1638c1d8
                  Next-hop reference count: 6
                  Next hop: #0 0.13.1.0.0.1 via demux0.1073753267, selected
                  Session Id: 0x0
                  State: <Active Int>
                  Age: 4:17
                  Validation State: unverified
                  Task: RPD Unix Domain Server./var/run/rpd_serv.local
                  Announcement bits (2): 0-KRT 4-Resolve tree 2
                  AS path: I
2001:db8:4641:1::/128 (1 entry, 1 announced)
  TSI:
  KRT in-kernel 2001:db8:4641:1::/128 -> {#0 0.13.1.0.0.1}
          *Access-internal Preference: 12
                  Next hop type: Router, Next hop index: 74548
                  Address: 0x1638c1d8
                  Next-hop reference count: 6
                  Next hop: #0 0.13.1.0.0.1 via demux0.1073753267, selected
                  Session Id: 0x0
                  State: <Active Int>
                  Age: 4:17
                  Validation State: unverified
                  Task: RPD Unix Domain Server./var/run/rpd_serv.local
                  Announcement bits (2): 0-KRT 4-Resolve tree 2
                  AS path: I
```

show route extensive (BGP PIC Edge)

```
user@host> show route 198.51.100.6 extensive
ed.inet.0: 6 destinations, 9 routes (6 active, 0 holddown, 0 hidden)
  198.51.100.6/32 (3 entries, 2 announced)
```

Copyright © 2018, Juniper Networks, Inc.

```
          State: <CalcForwarding>
  TSI:
  KRT in-kernel 198.51.100.6/32 -> {indirect(1048574), indirect(1048577)}
  Page 0 idx 0 Type 1 val 9219e30
      Nexthop: Self
      AS path: [2] 3 I
      Communities: target:2:1
  Path 198.51.100.6 from 198.51.100.4 Vector len 4.  Val: 0
..
          #Multipath Preference: 255
                  Next hop type: Indirect
                  Address: 0x93f4010
                  Next-hop reference count: 2
..
                  Protocol next hop: 198.51.1001.4
                  Push 299824
                  Indirect next hop: 944c000 1048574 INH Session ID: 0x3
                  Indirect next hop: weight 0x1
                  Protocol next hop: 198.51.100.5
                  Push 299824
                  Indirect next hop: 944c1d8 1048577 INH Session ID: 0x4
                  Indirect next hop: weight 0x4000
                  State: <ForwardingOnly Int Ext)>
                  Inactive reason: Forwarding use only
                  Age: 25         Metric2: 15
                  Validation State: unverified
                  Task: RT
                  Announcement bits (1): 0-KRT
                  AS path: 3 I
                  Communities: target:2:1
```

show route extensive (FRR and LFA)

```
user@host> show route 203.0.113.20 extensive
inet.0: 46 destinations, 49 routes (45 active, 0 holddown, 1 hidden)
203.0.113.20/24 (2 entries, 1 announced)
        State: FlashAll
TSI:
KRT in-kernel 203.0.113.20/24 -> {Push 299776, Push 299792}
        *RSVP   Preference: 7/1
                Next hop type: Router, Next hop index: 1048574
                Address: 0xbbbc010
                Next-hop reference count: 5
                Next hop: 203.0.113.112 via ge-2/1/8.0 weight 0x1, selected
                Label-switched-path europa-d-to-europa-e
                Label operation: Push 299776
                Label TTL action: prop-ttl
                Session Id: 0x201
                Next hop: 203.0.113.122 via ge-2/1/4.0 weight 0x4001
                Label-switched-path europa-d-to-europa-e
                Label operation: Push 299792
                Label TTL action: prop-ttl
                Session Id: 0x202
                State: Active Int
                Local AS:   64500
                Age: 5:31  Metric: 2
                Task: RSVP
                Announcement bits (1): 0-KRT
                AS path: I
         OSPF   Preference: 10
                Next hop type: Router, Next hop index: 615
```

 Copyright © 2018, Juniper Networks, Inc.

```
                Address: 0xb9d78c4
                Next-hop reference count: 7
                Next hop: 203.0.113.112 via ge-2/1/8.0, selected
                Session Id: 0x201
                State: Int
                Inactive reason: Route Preference
                Local AS:   64500
                Age: 5:35  Metric: 3
                Area: 0.0.0.0
                Task: OSPF
                AS path: I
```

show route extensive (IS-IS)

```
user@host> show route extensive
IS-IS Preference: 15
                Level: 1
                Next hop type: Router, Next hop index: 1048577
                Address: 0xXXXXXXXXXX
                Next-hop reference count: YY
                Next hop: 203.0.113.22 via ae1.0 balance 43%, selected
                Session Id: 0x141
                Next hop: 203.0.113.22 via ae0.0 balance 57%
```

show route extensive (Route Reflector)

```
user@host> show route extensive
203.0.113.0/8 (1 entry, 1 announced)

TSI:
KRT in-kernel 203.0.113.0/8 -> {indirect(40)}
        *BGP    Preference: 170/-101
                Source: 192.168.4.214
                Protocol next hop: 198.51.100.192 Indirect next hop: 84ac908 40
                State: <Active Int Ext>
                Local AS: 65548 Peer AS: 65548
                Age: 3:09       Metric: 0       Metric2: 0
                Task: BGP_65548.192.168.4.214+1033
                Announcement bits (2): 0-KRT 4-Resolve inet.0
                AS path: 65544 64507 I <Originator>
                Cluster list:  198.51.100.1
                Originator ID: 203.0.113.88
                Communities: 7777:7777
                Localpref: 100
                Router ID: 203.0.113.4
                Indirect next hops: 1
                        Protocol next hop: 203.0.113.192 Metric: 0
                        Indirect next hop: 84ac908 40
                        Indirect path forwarding next hops: 0
                                Next hop type: Discard
```

show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Next-hop reference count: 3
                Address: 0x9097d90
```

```
                Next hop: via vt-0/1/0.1
                Next-hop index: 661
                Label operation: Pop
                Address: 0x9172130
                Next hop: via so-0/0/3.0
                Next-hop index: 654
                Label operation: Swap 299872
                State: **Active Int>
                Local AS:  64511
                Age: 8:20       Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
              FECs bound to route: P2MP root-addr 203.0.113.166, grp 203.0.113.1,
 src 192.168.142.2
```

show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

user@host> **show route label 301568 detail**

```
mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Address: 0x2735208
                Next-hop reference count: 3
                Next hop type: Router, Next hop index: 1397
                Address: 0x2735d2c
                Next-hop reference count: 3
                Next hop: 203.0.113.82 via ge-1/2/22.0
                Label operation: Pop
                Load balance label:  None;
                Next hop type: Router, Next hop index: 1395
                Address: 0x2736290
                Next-hop reference count: 3
                Next hop: 203.0.113.2 via ge-1/2/18.0
                Label operation: Pop
                Load balance label:  None;
                State: <Active Int AckRequest MulticastRPF>
                Local AS:    64500
                Age: 54:05      Metric: 1
                Validation State: unverified
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
              FECs bound to route: P2MP root-addr 198.51.100.1, grp: 203.0.113.1,
 src: 192.168.219.11
                Primary Upstream : 198.51.100.3:0--198.51.100.2:0
                   RPF Nexthops  :
                        ge-1/2/15.0, 10.2.94.1, Label: 301568, weight: 0x1
                        ge-1/2/14.0, 10.2.3.1, Label: 301568, weight: 0x1
                Backup Upstream : 198.51.100.3:0--198.51.100.6:0
                   RPF Nexthops :
                       ge-1/2/20.0, 198.51.100.96, Label: 301584, weight: 0xfffe

                       ge-1/2/19.0, 198.51.100.36, Label: 301584, weight: 0xfffe
```

 Copyright © 2018, Juniper Networks, Inc.

show route flow validation

List of Syntax

Syntax on page 2265
Syntax (EX Series Switches) on page 2265

Syntax

show route flow validation
<brief | detail>
<*ip-prefix*>
<table *table-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route flow validation
<brief | detail>
<*ip-prefix*>
<table *table-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display flow route information.

Options

none—Display flow route information.

brief | detail—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.

ip-prefix—(Optional) IP address for the flow route.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

table *table-name*—(Optional) Display flow route information for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the **show route flow validation inet** command).

Required Privilege Level

view

List of Sample Output

show route flow validation on page 2266
show route flow validation (IPv6) on page 2266

Output Fields

Table 203 on page 2265 lists the output fields for the **show route flow validation** command. Output fields are listed in the approximate order in which they appear.

Table 203: show route flow validation Output Fields

Field Name	Field Description	Level of Output
routing-table-name	Name of the routing table (for example, inet.0).	All levels

Table 203: show route flow validation Output Fields (continued)

Field Name	Field Description	Level of Output
prefix	Route address.	All levels
Active unicast route	Active route in the routing table.	All levels
Dependent flow destinations	Number of flows for which there are routes in the routing table.	All levels
Origin	Source of the route flow.	All levels
Neighbor AS	Autonomous system identifier of the neighbor.	All levels
Flow destination	Number of entries and number of destinations that match the route flow.	All levels
Unicast best match	Destination that is the best match for the route flow.	All levels
Flags	Information about the route flow.	All levels

Sample Output

show route flow validation

```
user@host> show route flow validation
inet.0:
10.0.5.0/24Active unicast route
Dependent flow destinations: 1
Origin: 192.168.224.218, Neighbor AS: 64501
Flow destination (3 entries, 1 match origin)
Unicast best match: 10.0.5.0/24
Flags: SubtreeApex Consistent
```

show route flow validation (IPv6)

```
user@host> show route flow validation
inet6.0:
2001:db8::11:11:11:0/120
                Active unicast route
                    Dependent flow destinations: 2
                    Origin: 2001:db8::13:14:2:2, Neighbor AS: 2000
2001:db8::11:11:11:10/128
                Flow destination (1 entries, 1 match origin, next-as)
                    Unicast best match: 2001:db8::11:11:11:0/120
                    Flags: Consistent
2001:db8::11:11:11:30/128
                Flow destination (1 entries, 1 match origin, next-as)
                    Unicast best match: 2001:db8::11:11:11:0/120
                    Flags: Consistent
```


Copyright © 2018, Juniper Networks, Inc.

show route forwarding-table

List of Syntax

Syntax on page 2267
Syntax (MX Series Routers) on page 2267
Syntax (TX Matrix and TX Matrix Plus Routers) on page 2267

Syntax

show route forwarding-table
<detail | extensive | summary>
<all>
<ccc *interface-name*>
<destination *destination-prefix*>
<family *family* | matching *matching*>
<interface-name *interface-name*>
<label *name*>
<matching *matching*>
<multicast>
<table (default | *logical-system-name/routing-instance-name* | *routing-instance-name*)>
<vlan (all | *vlan-name*)>
<vpn *vpn*>

Syntax (MX Series Routers)

show route forwarding-table
<detail | extensive | summary>
<all>
<bridge-domain (all | *domain-name*)>
<ccc *interface-name*>
<destination *destination-prefix*>
<family *family* | matching *matching*>
<interface-name *interface-name*>
<label *name*>
<learning-vlan-id *learning-vlan-id*>
<matching *matching*>
<multicast>
<table (default | *logical-system-name/routing-instance-name* | *routing-instance-name*)>
<vlan (all | *vlan-name*)>
<vpn *vpn*>

Syntax (TX Matrix and TX Matrix Plus Routers)

show route forwarding-table
<detail | extensive | summary>
<all>
<ccc *interface-name*>
<destination *destination-prefix*>
<family *family* | matching *matching*>
<interface-name *interface-name*>
<matching *matching*>
<label *name*>
<lcc *number*>
<multicast>
<table *routing-instance-name*>
<vpn *vpn*>

Release Information

Command introduced before Junos OS Release 7.4.
Option **bridge-domain** introduced in Junos OS Release 7.5

Copyright © 2018, Juniper Networks, Inc.

Option **learning-vlan-id** introduced in Junos OS Release 8.4
Options **all** and **vlan** introduced in Junos OS Release 9.6.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Display the Routing Engine's forwarding table, including the network-layer prefixes and their next hops. This command is used to help verify that the routing protocol process has relayed the correction information to the forwarding table. The Routing Engine constructs and maintains one or more routing tables. From the routing tables, the Routing Engine derives a table of active routes, called the forwarding table.

NOTE: The Routing Engine copies the forwarding table to the Packet Forwarding Engine, the part of the router that is responsible for forwarding packets. To display the entries in the Packet Forwarding Engine's forwarding table, use the **show pfe route** command.

Options **none**—Display the routes in the forwarding tables. By default, the **show route forwarding-table** command does not display information about private, or internal, forwarding tables.

detail | extensive | summary—(Optional) Display the specified level of output.

all—(Optional) Display routing table entries for all forwarding tables, including private, or internal, tables.

bridge-domain (all | ***bridge-domain-name*****)**—(MX Series routers only) (Optional) Display route entries for all bridge domains or the specified bridge domain.

ccc ***interface-name***—(Optional) Display route entries for the specified circuit cross-connect interface.

destination ***destination-prefix***—(Optional) Destination prefix.

family ***family***—(Optional) Display routing table entries for the specified family: **bridge (ccc | destination | detail | extensive | interface-name | label | learning-vlan-id | matching | multicast | summary | table | vlan | vpn)**, **ethernet-switching**, **evpn**, **fibre-channel**, **fmembers**, **inet**, **inet6**, **iso**, **mcsnoop-inet**, **mcsnoop-inet6**, **mpls**, **satellite-inet**, **satellite-inet6**, **satellite-vpls**, **tnp**, **unix**, **vpls**, or **vlan-classification**.

interface-name ***interface-name***—(Optional) Display routing table entries for the specified interface.

label ***name***—(Optional) Display route entries for the specified label.

lcc ***number***—(TX Matrix and TX matrix Plus routers only) (Optional) On a routing matrix composed of a TX Matrix router and T640 routers, display information for the specified T640 router (or line-card chassis) connected to the TX Matrix router. On a routing matrix composed of the TX Matrix Plus router and T1600 or T4000 routers,

 Copyright © 2018, Juniper Networks, Inc.

display information for the specified router (line-card chassis) connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

learning-vlan-id *learning-vlan-id*—(MX Series routers only) (Optional) Display learned information for all VLANs or for the specified VLAN.

matching *matching*—(Optional) Display routing table entries matching the specified prefix or prefix length.

multicast—(Optional) Display routing table entries for multicast routes.

table —(Optional) Display route entries for all the routing tables in the main routing instance or for the specified routing instance. If your device supports logical systems, you can also display route entries for the specified logical system and routing instance. To view the routing instances on your device, use the show route instance command.

vlan (all | *vlan-name*)—(Optional) Display information for all VLANs or for the specified VLAN.

vpn *vpn*—(Optional) Display routing table entries for a specified VPN.

Required Privilege Level view

List of Sample Output

Copyright © 2018, Juniper Networks, Inc.

show route forwarding-table table default on page 2281
show route forwarding-table table
logical-system-name/routing-instance-name on page 2282
show route forwarding-table vpn on page 2283

Output Fields

Table 204 on page 2270 lists the output fields for the **show route forwarding-table** command. Output fields are listed in the approximate order in which they appear. Field names might be abbreviated (as shown in parentheses) when no level of output is specified, or when the **detail** keyword is used instead of the **extensive** keyword.

Table 204: show route forwarding-table Output Fields

Field Name	Field Description	Level of Output
Logical system	Name of the logical system. This field is displayed if you specify the **table** ***logical-system-name/routing-instance-name*** option on a device that is configured for and supports logical systems.	All levels
Routing table	Name of the routing table (for example, inet, inet6, mpls).	All levels

Copyright © 2018, Juniper Networks, Inc.

Table 204: show route forwarding-table Output Fields (continued)

Field Name	Field Description	Level of Output
Enabled protocols	The features and protocols that have been enabled for a given routing table. This field can contain the following values: • BUM hashing—BUM hashing is enabled. • MAC Stats—Mac Statistics is enabled. • Bridging—Routing instance is a normal layer 2 bridge. • No VLAN—No VLANs are associated with the bridge domain. • All VLANs—The **vlan-id all** statement has been enabled for this bridge domain. • Single VLAN—Single VLAN ID is associated with the bridge domain. • MAC action drop—New MACs will be dropped when the MAC address limit is reached. • Dual VLAN—Dual VLAN tags are associated with the bridge domain • No local switching—No local switching is enabled for this routing instance.. • Learning disabled—Layer 2 learning is disabled for this routing instance. • MAC limit reached—The maximum number of MAC addresses that was configured for this routing instance has been reached. • VPLS—The VPLS protocol is enabled. • No IRB l2-copy—The no-irb-layer-2-copy feature is enabled for this routing instance. • ACKed by all peers—All peers have acknowledged this routing instance. • BUM Pruning—BUM pruning is enabled on the VPLS instance. • Def BD VXLAN—VXLAN is enabled for the default bridge domain. • EVPN—EVPN protocol is enabled for this routing instance. • Def BD OVSDB—Open vSwitch Database (OVSDB) is enabled on the default bridge domain. • Def BD Ingress replication—VXLAN ingress node replication is enabled on the default bridge domain. • L2 backhaul—Layer 2 backhaul is enabled. • FRR optimize—Fast reroute optimization • MAC pinning—MAC pinning is enabled for this bridge domain. • MAC Aging Timer—The MAC table aging time is set per routing instance. • EVPN VXLAN—This routing instance supports EVPN with VXLAN encapsulation. • PBBN—This routing instance is configured as a provider backbone bridged network. • PBN—This routing instance is configured as a provider bridge network. • ETREE—The ETREE protocol is enabled on this EVPN routing instance. • ARP/NDP suppression—EVPN ARP NDP suppression is enabled in this routing instance. • Def BD EVPN VXLAN—EVPN VXLAN is enabled for the default bridge domain. • MPLS control word—Control word is enabled for this MPLS routing instance.	All levels
Address family	Address family (for example, **IP**, **IPv6**, **ISO**, **MPLS**, and **VPLS**).	All levels
Destination	Destination of the route.	**detail extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 204: show route forwarding-table Output Fields (continued)

Field Name	Field Description	Level of Output
Route Type (Type)	How the route was placed into the forwarding table. When the **detail** keyword is used, the route type might be abbreviated (as shown in parentheses): • **cloned** (**clon**)—(TCP or multicast only) Cloned route. • **destination** (**dest**)—Remote addresses directly reachable through an interface. • **destination down** (**iddn**)—Destination route for which the interface is unreachable. • **interface cloned** (**ifcl**)—Cloned route for which the interface is unreachable. • **route down** (**ifdn**)—Interface route for which the interface is unreachable. • **ignore** (**ignr**)—Ignore this route. • **interface** (**intf**)—Installed as a result of configuring an interface. • **permanent** (**perm**)—Routes installed by the kernel when the routing table is initialized. • **user**—Routes installed by the routing protocol process or as a result of the configuration.	All levels
Route Reference (RtRef)	Number of routes to reference.	**detail extensive**
Flags	Route type flags: • **none**—No flags are enabled. • **accounting**—Route has accounting enabled. • **cached**—Cache route. • **incoming-iface** ***interface-number***—Check against incoming interface. • **prefix load balance**—Load balancing is enabled for this prefix. • **rt nh decoupled**—Route has been decoupled from the next hop to the destination. • **sent to PFE**—Route has been sent to the Packet Forwarding Engine. • **static**—Static route.	**extensive**
Next hop	IP address of the next hop to the destination.	**detail extensive**

 Copyright © 2018, Juniper Networks, Inc.

Table 204: show route forwarding-table Output Fields (continued)

Field Name	Field Description	Level of Output
Next hop Type (Type)	Next-hop type. When the **detail** keyword is used, the next-hop type might be abbreviated (as indicated in parentheses): • **broadcast (bcst)**—Broadcast. • **deny**—Deny. • **discard (dscd)** —Discard. • **hold**—Next hop is waiting to be resolved into a unicast or multicast type. • **indexed (idxd)**—Indexed next hop. • **indirect (indr)**—Indirect next hop. • **local (locl)**—Local address on an interface. • **routed multicast (mcrt)**—Regular multicast next hop. • **multicast (mcst)**—Wire multicast next hop (limited to the LAN). • **multicast discard (mdsc)**—Multicast discard. • **multicast group (mgrp)**—Multicast group member. • **receive (recv)**—Receive. • **reject (rjct)**—Discard. An ICMP unreachable message was sent. • **resolve (rslv)**—Resolving the next hop. • **unicast (ucst)**—Unicast. • **unilist (ulst)**—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.	**detail extensive**
Index	Software index of the next hop that is used to route the traffic for a given prefix.	**detail extensive** none
Route interface-index	Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned.	**extensive**
Reference (NhRef)	Number of routes that refer to this next hop.	**detail extensive** none
Next-hop interface (Netif)	Interface used to reach the next hop.	**detail extensive** none
Weight	Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible (see the **Balance** field description).	**extensive**
Balance	Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a router is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.	**extensive**
RPF interface	List of interfaces from which the prefix can be accepted. Reverse path forwarding (RPF) information is displayed only when **rpf-check** is configured on the interface.	**extensive**

Sample Output

show route forwarding-table

```
user@host> show route forwarding-table
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    46     4
0.0.0.0/32         perm     0                    dscd    44     1
172.16.1.0/24         ifdn     0                    rslv   608     1 ge-2/0/1.0
172.16.1.0/32        iddn     0 172.16.1.0           recv   606     1 ge-2/0/1.0
172.16.1.1/32         user     0                    rjct    46     4
172.16.1.1/32         intf     0 172.16.1.1            locl   607     2
172.16.1.1/32         iddn     0 172.16.1.1            locl   607     2
172.16.1.255/32       iddn     0 ff:ff:ff:ff:ff:ff  bcst   605     1 ge-2/0/1.0
10.0.0.0/24        intf     0                    rslv   616     1 ge-2/0/0.0
10.0.0.0/32        dest     0 10.0.0.0           recv   614     1 ge-2/0/0.0
10.0.0.1/32        intf     0 10.0.0.1           locl   615     2
10.0.0.1/32        dest     0 10.0.0.1           locl   615     2
10.0.0.255/32      dest     0 10.0.0.255         bcst   613     1 ge-2/0/0.0
10.1.1.0/24        ifdn     0                    rslv   612     1 ge-2/0/1.0
10.1.1.0/32        iddn     0 10.1.1.0           recv   610     1 ge-2/0/1.0
10.1.1.1/32        user     0                    rjct    46     4
10.1.1.1/32        intf     0 10.1.1.1           locl   611     2
10.1.1.1/32        iddn     0 10.1.1.1           locl   611     2
10.1.1.255/32      iddn     0 ff:ff:ff:ff:ff:ff  bcst   609     1 ge-2/0/1.0
10.206.0.0/16      user     0 10.209.63.254      ucst   419    20 fxp0.0
10.209.0.0/16      user     1 0:12:1e:ca:98:0    ucst   419    20 fxp0.0
10.209.0.0/18      intf     0                    rslv   418     1 fxp0.0
10.209.0.0/32      dest     0 10.209.0.0         recv   416     1 fxp0.0
10.209.2.131/32    intf     0 10.209.2.131       locl   417     2
10.209.2.131/32    dest     0 10.209.2.131       locl   417     2
10.209.17.55/32    dest     0 0:30:48:5b:78:d2   ucst   435     1 fxp0.0
10.209.63.42/32    dest     0 0:23:7d:58:92:ca   ucst   434     1 fxp0.0
10.209.63.254/32   dest     0 0:12:1e:ca:98:0    ucst   419    20 fxp0.0
10.209.63.255/32   dest     0 10.209.63.255      bcst   415     1 fxp0.0
10.227.0.0/16      user     0 10.209.63.254      ucst   419    20 fxp0.0

...

Routing table: iso
ISO:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                   rjct    27     1
47.0005.80ff.f800.0000.0108.0003.0102.5524.5220.00
intf     0                 locl     28     1

Routing table: inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                   rjct     6     1
ff00::/8           perm     0                   mdsc     4     1
ff02::1/128        perm     0 ff02::1           mcst     3     1

Routing table: ccc
MPLS:
Interface.Label    Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct 16     1
100004(top)fe-0/0/1.0
```

 Copyright © 2018, Juniper Networks, Inc.

show route forwarding-table detail

```
user@host> show route forwarding-table detail
Routing table: inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            user     2 0:90:69:8e:b1:1b   ucst   132     4 fxp0.0
default            perm     0                    rjct    14     1
10.1.1.0/24        intf     0 ff.3.0.21          ucst   322     1 so-5/3/0.0
10.1.1.0/32        dest     0 10.1.1.0           recv   324     1 so-5/3/0.0
10.1.1.1/32        intf     0 10.1.1.1           locl   321     1
10.1.1.255/32      dest     0 10.1.1.255         bcst   323     1 so-5/3/0.0
10.21.21.0/24      intf     0 ff.3.0.21          ucst   326     1 so-5/3/0.0
10.21.21.0/32      dest     0 10.21.21.0         recv   328     1 so-5/3/0.0
10.21.21.1/32      intf     0 10.21.21.1         locl   325     1
10.21.21.255/32    dest     0 10.21.21.255       bcst   327     1 so-5/3/0.0
127.0.0.1/32       intf     0 127.0.0.1          locl   320     1
172.17.28.19/32    clon     1 192.168.4.254      ucst   132     4 fxp0.0
172.17.28.44/32    clon     1 192.168.4.254      ucst   132     4 fxp0.0

...

Routing table: private1__.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    46     1
10.0.0.0/8         intf     0                    rslv   136     1 fxp1.0
10.0.0.0/32        dest     0 10.0.0.0           recv   134     1 fxp1.0
10.0.0.4/32        intf     0 10.0.0.4           locl   135     2
10.0.0.4/32        dest     0 10.0.0.4           locl   135     2

...

Routing table: iso
ISO:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    38     1

Routing table: inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    22     1
ff00::/8           perm     0                    mdsc    21     1
ff02::1/128        perm     0 ff02::1            mcst    17     1

...

Routing table: mpls
MPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct 28     1
```

show route forwarding-table destination extensive (Weights and Balances)

```
user@host> show route forwarding-table destination 3.4.2.1 extensive
Routing table: inet [Index 0]
Internet:

Destination:  3.4.2.1/32
  Route type: user
```

```
  Route reference: 0                   Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: unilist               Index: 262143   Reference: 1
  Nexthop: 172.16.4.4
  Next-hop type: unicast               Index: 335      Reference: 2
  Next-hop interface: so-1/1/0.0       Weight: 22      Balance: 3
  Nexthop: 145.12.1.2
  Next-hop type: unicast               Index: 337      Reference: 2
  Next-hop interface: so-0/1/2.0       Weight: 33      Balance: 33
```

show route forwarding-table extensive

```
user@host> show route forwarding-table extensive
Routing table: inet [Index 0]
Internet:

Destination:  default
  Route type: user
  Route reference: 2                   Route interface-index: 0
  Flags: sent to PFE
  Nexthop: 00:00:5E:00:53:1b
  Next-hop type: unicast               Index: 132      Reference: 4
  Next-hop interface: fxp0.0

Destination:  default
  Route type: permanent
  Route reference: 0                   Route interface-index: 0
  Flags: none
  Next-hop type: reject                Index: 14       Reference: 1

Destination:  127.0.0.1/32
  Route type: interface
  Route reference: 0                   Route interface-index: 0
  Flags: sent to PFE
  Nexthop: 127.0.0.1
  Next-hop type: local                 Index: 320      Reference: 1

...

Routing table: private1__.inet [Index 1]
Internet:

Destination:  default
  Route type: permanent
  Route reference: 0                   Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: reject                Index: 46       Reference: 1

Destination:  10.0.0.0/8
  Route type: interface
  Route reference: 0                   Route interface-index: 3
  Flags: sent to PFE
  Next-hop type: resolve               Index: 136      Reference: 1
  Next-hop interface: fxp1.0

...

Routing table: iso [Index 0]
ISO:

Destination:  default
```


Copyright © 2018, Juniper Networks, Inc.

```
  Route type: permanent
  Route reference: 0                   Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: reject                Index: 38       Reference: 1

Routing table: inet6 [Index 0]
Internet6:

Destination:  default
  Route type: permanent
  Route reference: 0                   Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: reject                Index: 22       Reference: 1

Destination:  ff00::/8
  Route type: permanent
  Route reference: 0                   Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: multicast discard     Index: 21       Reference: 1

...

Routing table: private1__.inet6 [Index 1]
Internet6:

Destination:  default
  Route type: permanent
  Route reference: 0                   Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: reject                Index: 54       Reference: 1

Destination:  fe80::2a0:a5ff:fe3d:375/128
  Route type: interface
  Route reference: 0                   Route interface-index: 0
  Flags: sent to PFE
  Nexthop: fe80::2a0:a5ff:fe3d:375
  Next-hop type: local                 Index: 75       Reference: 1

...
```

show route forwarding-table extensive (RPF)

The next example is based on the following configuration, which enables an RPF check on all routes that are learned from this interface, including the interface route:

```
so-1/1/0 {
    unit 0 {
        family inet {
            rpf-check;
            address 192.0.2.2/30;
        }
    }
}
```

```
user@host> show route forwarding-table extensive
Routing table: inet [Index 0]
Internet:
...
...
Destination:  192.0.2.3/32
```

```
  Route type: destination
  Route reference: 0                   Route interface-index: 67
  Flags: sent to PFE
  Nexthop: 192.0.2.3
  Next-hop type: broadcast             Index: 328      Reference: 1
  Next-hop interface: so-1/1/0.0
  RPF interface: so-1/1/0.0
```

show route forwarding-table (dynamic list next hop)

The **show route forwarding table** output shows the two next hop elements for a multihomed EVPN destination.

```
user@host> show route forwarding-table label 299952 extensive
MPLS:

Destination:  299952
  Route type: user
  Route reference: 0                   Route interface-index: 0
  Multicast RPF nh index: 0
  P2mpidx: 0
  Flags: sent to PFE, rt nh decoupled
  Next-hop type: indirect              Index: 1048575  Reference: 2
  Nexthop:
  Next-hop type: composite             Index: 601      Reference: 2
  Next-hop type: indirect              Index: 1048574  Reference: 3
  Nexthop: 1.0.0.4
  Next-hop type: Push 301632, Push 299776(top) Index: 600 Reference: 2
  Load Balance Label: None
  Next-hop interface: ge-0/0/1.0
  Next-hop type: indirect              Index: 1048577  Reference: 3
  Nexthop: 1.0.0.4
  Next-hop type: Push 301344, Push 299792(top) Index: 619 Reference: 2
  Load Balance Label: None
  Next-hop interface: ge-0/0/1.0
```

After one of the PE router has been disabled in the EVPN multihomed network, the same **show route forwarding table** output command shows one next hop element and one empty next hop element.

```
user@host> show route forwarding-table label 299952 extensive
Routing table: default.mpls [Index 0]
MPLS:

Destination:  299952
  Route type: user
  Route reference: 0                   Route interface-index: 0
  Multicast RPF nh index: 0
  P2mpidx: 0
  Flags: sent to PFE, rt nh decoupled
  Next-hop type: indirect              Index: 1048575  Reference: 2
  Nexthop:
  Next-hop type: composite             Index: 601      Reference: 2
  Next-hop type: indirect              Index: 1048577  Reference: 3
  Nexthop: 1.0.0.4
  Next-hop type: Push 301344, Push 299792(top) Index: 619 Reference: 2
  Load Balance Label: None
  Next-hop interface: ge-0/0/1.0
```

 Copyright © 2018, Juniper Networks, Inc.

show route forwarding-table family mpls

```
user@host> show route forwarding-table family mpls
Routing table: mpls
MPLS:
Destination        Type RtRef Next hop          Type Index NhRef Netif
default            perm     0                   rjct    19     1
0                  user     0                   recv    18     3
1                  user     0                   recv    18     3
2                  user     0                   recv    18     3
100000             user     0 10.31.1.6         swap  100001     fe-1/1/0.0
800002             user     0                   Pop                vt-0/3/0.32770

vt-0/3/0.32770 (VPLS)
                   user     0                   indr   351     4
                                                Push 800000, Push 100002(top)
so-0/0/0.0
```

show route forwarding-table family mpls ccc ge-0/0/1.1004

```
user@host>show route forwarding-table mpls ccc ge-0/0/1.1004
Routing table: default.mpls
MPLS:
Destination        Type RtRef Next hop          Type  Index    NhRef Netif
ge-0/0/1.1004   (CCC) user     0                    ulst   1048577     2
                                                    comp      754         3
                                                   comp      755         3
                                                comp      756         3

Routing table: __mpls-oam__.mpls
MPLS:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
default            perm     0                    dscd      556     1
```

show route forwarding-table family vpls

```
user@host> show route forwarding-table family vpls
Routing table: green.vpls
VPLS:
Destination        Type RtRef Next hop          Type Index NhRef Netif
default            dynm     0                   flood   353     1
default            perm     0                   rjct    298     1
fe-0/1/0.0         dynm     0                   flood   355     1
00:00:5E:00:53:1f/48               <<<<<Remote CE

                   dynm     0                   indr   351     4
                                                Push 800000, Push 100002(top)
so-0/0/0.0
00:00:5E:00:53:1f/48               <<<<<<Local CE

                   dynm     0                   ucst   354     2 fe-0/1/0.0
```

show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled)

```
user@host> show route forwarding-table vpls
Routing table: green.vpls
VPLS:
Enabled protocols: BUM hashing
Destination        Type RtRef Next hop           Type Index    NhRef Netif
```

```
default            perm     0                     dscd      519     1
lsi.1048832        intf     0                     indr  1048574     4
                              172.16.3.2            Push 262145      621     2
ge-3/0/0.0
00:00:5E:00:53:01/48 user     0                   ucst      590     5 ge-2/3/9.0
0x30003/51         user     0                     comp      627     2
ge-2/3/9.0         intf     0                     ucst      590     5 ge-2/3/9.0
ge-3/1/3.0         intf     0                     ucst      619     4 ge-3/1/3.0
0x30002/51         user     0                     comp      600     2
0x30001/51         user     0                     comp      597     2
```

show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled with MAC Statistics)

```
user@host> show route forwarding-table vpls
Routing table: green.vpls
VPLS:
Enabled protocols: BUM hashing, MAC Stats
Destination        Type RtRef Next hop           Type Index    NhRef Netif
default            perm     0                     dscd      519     1
lsi.1048834        intf     0                     indr  1048574     4
                              172.16.3.2            Push 262145      592     2
ge-3/0/0.0
00:19:e2:25:d0:01/48 user     0                   ucst      590     5 ge-2/3/9.0
0x30003/51         user     0                     comp      630     2
ge-2/3/9.0         intf     0                     ucst      590     5 ge-2/3/9.0
ge-3/1/3.0         intf     0                     ucst      591     4 ge-3/1/3.0
0x30002/51         user     0                     comp      627     2
0x30001/51         user     0                     comp      624     2
```

show route forwarding-table family vpls extensive

```
user@host> show route forwarding-table family vpls extensive
Routing table: green.vpls [Index 2]
VPLS:

Destination:  default
  Route type: dynamic
  Route reference: 0                   Route interface-index: 72
  Flags: sent to PFE
  Next-hop type: flood                 Index: 289      Reference: 1
  Next-hop type: unicast               Index: 291      Reference: 3
  Next-hop interface: fe-0/1/3.0
  Next-hop type: unicast               Index: 290      Reference: 3
  Next-hop interface: fe-0/1/2.0

Destination:  default
  Route type: permanent
  Route reference: 0                   Route interface-index: 0
  Flags: none
  Next-hop type: discard               Index: 341      Reference: 1

Destination:  fe-0/1/2.0
  Route type: dynamic
  Route reference: 0                   Route interface-index: 69
  Flags: sent to PFE
  Next-hop type: flood                 Index: 293      Reference: 1
  Next-hop type: indirect              Index: 363      Reference: 4
  Next-hop type: Push 800016
  Next-hop interface: at-1/0/1.0
```

 Copyright © 2018, Juniper Networks, Inc.

```
  Next-hop type: indirect              Index: 301      Reference: 5
  Next hop: 10.31.3.2
  Next-hop type: Push 800000
  Next-hop interface: fe-0/1/1.0
  Next-hop type: unicast               Index: 291      Reference: 3
  Next-hop interface: fe-0/1/3.0

Destination:  fe-0/1/3.0
  Route type: dynamic
  Route reference: 0                   Route interface-index: 70
  Flags: sent to PFE
  Next-hop type: flood                 Index: 292      Reference: 1
  Next-hop type: indirect              Index: 363      Reference: 4
  Next-hop type: Push 800016
  Next-hop interface: at-1/0/1.0
  Next-hop type: indirect              Index: 301      Reference: 5
  Next hop: 10.31.3.2
  Next-hop type: Push 800000
  Next-hop interface: fe-0/1/1.0
  Next-hop type: unicast               Index: 290      Reference: 3
  Next-hop interface: fe-0/1/2.0

Destination:  00:00:5E:00:53:01/48
  Route type: dynamic
  Route reference: 0                   Route interface-index: 70
  Flags: sent to PFE, prefix load balance
  Next-hop type: unicast               Index: 291      Reference: 3
  Next-hop interface: fe-0/1/3.0
  Route used as destination:
    Packet count:         6640    Byte count:      675786
  Route used as source
    Packet count:         6894    Byte count:      696424

Destination:  00:00:5E:00:53:04/48
  Route type: dynamic
  Route reference: 0                   Route interface-index: 69
  Flags: sent to PFE, prefix load balance
  Next-hop type: unicast               Index: 290      Reference: 3
  Next-hop interface: fe-0/1/2.0
  Route used as destination:
    Packet count:           96    Byte count:        8079
  Route used as source:
    Packet count:          296    Byte count:       24955

Destination:  00:00:5E:00:53:05/48
  Route type: dynamic
  Route reference: 0                   Route interface-index: 74
  Flags: sent to PFE, prefix load balance
  Next-hop type: indirect              Index: 301      Reference: 5
  Next hop: 10.31.3.2
  Next-hop type: Push 800000
  Next-hop interface: fe-0/1/1.0
```

show route forwarding-table table default

```
user@host> show route forwarding-table table default
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    36     2
0.0.0.0/32         perm     0                    dscd    34     1
```

```
10.0.60.0/30       user     0 10.0.60.13        ucst   713      5 fe-0/1/3.0
10.0.60.12/30      intf     0                   rslv   688      1 fe-0/1/3.0
10.0.60.12/32      dest     0 10.0.60.12        recv   686      1 fe-0/1/3.0
10.0.60.13/32      dest     0 0:5:85:8b:bc:22   ucst   713      5 fe-0/1/3.0
10.0.60.14/32      intf     0 10.0.60.14        locl   687      2
10.0.60.14/32      dest     0 10.0.60.14        locl   687      2
10.0.60.15/32      dest     0 10.0.60.15        bcst   685      1 fe-0/1/3.0
10.0.67.12/30      user     0 10.0.60.13        ucst   713      5 fe-0/1/3.0
10.0.80.0/30       ifdn     0 ff.3.0.21         ucst   676      1 so-0/0/1.0
10.0.80.0/32       dest     0 10.0.80.0         recv   678      1 so-0/0/1.0
10.0.80.2/32       user     0                   rjct    36      2
10.0.80.2/32       intf     0 10.0.80.2         locl   675      1
10.0.80.3/32       dest     0 10.0.80.3         bcst   677      1 so-0/0/1.0
10.0.90.12/30      intf     0                   rslv   684      1 fe-0/1/0.0
10.0.90.12/32      dest     0 10.0.90.12        recv   682      1 fe-0/1/0.0
10.0.90.14/32      intf     0 10.0.90.14        locl   683      2
10.0.90.14/32      dest     0 10.0.90.14        locl   683      2
10.0.90.15/32      dest     0 10.0.90.15        bcst   681      1 fe-0/1/0.0
10.5.0.0/16        user     0 192.168.187.126   ucst   324     15 fxp0.0
10.10.0.0/16       user     0 192.168.187.126   ucst   324     15 fxp0.0
10.13.10.0/23      user     0 192.168.187.126   ucst   324     15 fxp0.0
10.84.0.0/16       user     0 192.168.187.126   ucst   324     15 fxp0.0
10.150.0.0/16      user     0 192.168.187.126   ucst   324     15 fxp0.0
10.157.64.0/19     user     0 192.168.187.126   ucst   324     15 fxp0.0
10.209.0.0/16      user     0 192.168.187.126   ucst   324     15 fxp0.0

...

Routing table: default.iso
ISO:
Destination        Type RtRef Next hop          Type Index NhRef Netif
default            perm     0                   rjct    60     1

Routing table: default.inet6
Internet6:
Destination        Type RtRef Next hop          Type Index NhRef Netif
default            perm     0                   rjct    44     1
::/128             perm     0                   dscd    42     1
ff00::/8           perm     0                   mdsc    43     1
ff02::1/128        perm     0 ff02::1           mcst    39     1

Routing table: default.mpls
MPLS:
Destination        Type RtRef Next hop          Type Index NhRef Netif
default            perm     0                   dscd    50     1
```

show route forwarding-table table logical-system-name/routing-instance-name

```
user@host> show route forwarding-table table R4/vpn-red
Logical system: R4
Routing table: vpn-red.inet
Internet:
Destination        Type RtRef Next hop          Type Index NhRef Netif
default            perm     0                   rjct   563     1
0.0.0.0/32         perm     0                   dscd   561     2
172.16.0.1/32         user     0                   dscd   561     2
172.16.2.0/24         intf     0                   rslv   771     1 ge-1/2/0.3
172.16.2.0/32        dest     0 172.16.2.0          recv   769     1 ge-1/2/0.3
172.16.2.1/32         intf     0 172.16.2.1          locl   770     2
172.16.2.1/32         dest     0 172.16.2.1          locl   770     2
172.16.2.2/32         dest     0 0.4.80.3.0.1b.c0.d5.e4.bd.0.1b.c0.d5.e4.bc.8.0
```

 Copyright © 2018, Juniper Networks, Inc.

```
                                                 ucst   789      1 ge-1/2/0.3
172.16.2.255/32       dest     0 172.16.2.255        bcst   768      1 ge-1/2/0.3
172.16.233.0/4        perm     1                    mdsc   562      1
172.16.233.1/32       perm     0 172.16.233.1          mcst   558      1
255.255.255.255/32 perm     0                    bcst   559      1

Logical system: R4
Routing table: vpn-red.iso
ISO:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   608      1

Logical system: R4
Routing table: vpn-red.inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   708      1
::/128             perm     0                    dscd   706      1
ff00::/8           perm     0                    mdsc   707      1
ff02::1/128        perm     0 ff02::1            mcst   704      1

Logical system: R4
Routing table: vpn-red.mpls
MPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    dscd   638
```

show route forwarding-table vpn

```
user@host> show route forwarding-table vpn VPN-A
      Routing table:: VPN-A.inet
Internet:
      Destination        Type RtRef Nexthop           Type Index NhRef Netif
      default            perm     0                   rjct     4     4
      10.39.10.20/30     intf     0 ff.3.0.21         ucst    40     1
      so-0/0/0.0
      10.39.10.21/32     intf     0 10.39.10.21       locl    36     1
      10.255.14.172/32   user     0                   ucst    69     2
      so-0/0/0.0
      10.255.14.175/32   user     0                   indr    81     3
                                                      Push 100004, Push
      100004(top) so-1/0/0.0
      172.16.233.0/4        perm     2                   mdsc     5     3
      172.16.233.1/32       perm     0 172.16.233.1         mcst     1     8
      172.16.233.5/32       user     1 172.16.233.5         mcst     1     8
      255.255.255.255/32 perm     0                   bcst     2     3
```

On QFX5200, the results for this command look like this:

```
show route forwarding-table family mpls
Routing table: default.mpls
MPLS:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 dscd 65 1
0 user 0 recv 64 4
1 user 0 recv 64 4
2 user 0 recv 64 4
13 user 0 recv 64 4
300384 user 0 9.1.1.1 Pop 1711 2 xe-0/0/34.0
300384(S=0) user 0 9.1.1.1 Pop 1712 2 xe-0/0/34.0
```

```
300400 user 0 ulst 131071 2
                              10.1.1.2 Pop 1713 1 xe-0/0/38.0
                              172.16.11.2 Pop 1714 1 xe-0/0/40.0
300400(S=0) user 0 ulst 131072 2
                              10.1.1.2 Pop 1715 1 xe-0/0/38.0
                              172.16.11.2 Pop 1716 1 xe-0/0/40.0

Routing table: __mpls-oam__.mpls
MPLS:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 dscd 1681 1
```

 Copyright © 2018, Juniper Networks, Inc.

show route forwarding-table interface-name

Syntax

show route forwarding-table interface-name *interface-name*
<detail | extensive>
<all>

Release Information

Command introduced in Junos OS Release 9.6.

Description

Display the interfaces in the Routing Engine's forwarding table.

Options

none—Display information for the specified interface.

detail | extensive—(Optional) Display the specified level of output.

all—(Optional) Display all interfaces in the routing table.

Required Privilege Level

view

List of Sample Output

show route forwarding-table interface-name fe-0/1/1 on page 2286
show route forwarding-table interface-name all on page 2286
show route forwarding-table interface-name all detail on page 2287

Output Fields

Table 205 on page 2285 lists the output fields for the **show route forwarding-table interface-name** command. Output fields are listed in the approximate order in which they appear.

Table 205: show route forwarding-table interface-name Output Fields

Field Name	Field Description	Level of Output
Name	Name of the interface (for example **fe-0/1/1**, **lo0**, **ae0**, and so on).	All levels
MTU	Interface’s maximum transmission unit (MTU).	All levels
Afam	Configured address family (for example **inet**, **tnp**, **inet6**, and so on).	**detail extensive**
Network	Network information: • **<Link>**—Physical interface, not a logical interface. • **<PtoP>**—Point-to-point network. • **ipaddress**—Network address.	All levels
Address	Address of the interface. The address can be a MAC address, IPv4 address, IPv6 address, and so on.	All levels
IPkts	Number of packets received on the interface.	All levels
Ierr	Number of packets received on the interface with errors.	All levels

Copyright © 2018, Juniper Networks, Inc.

Table 205: show route forwarding-table interface-name Output Fields (continued)

Field Name	Field Description	Level of Output
Opkts	Number of packets transmitted or sent from the interface.	All levels
Oerr	Number of packets transmitted or sent from the interface with errors.	All levels
Coll	Number of packets that experienced collisions on the interface.	All levels

Sample Output

show route forwarding-table interface-name fe-0/1/1

```
user@host> show route forwarding-table interface-name fe-0/1/1
Name            Mtu Network       Address             Ipkts Ierr    Opkts Oerr   Coll
fe-0/1/1       1514 <Link>      00.05.85.88.cc.20         0    0        0    0      0
```

show route forwarding-table interface-name all

```
user@host> show route forwarding-table interface-name all
Name            Mtu Network       Address             Ipkts Ierr    Opkts Oerr   Coll
fxp0           1514 <Link>      00.a0.a5.56.03.83    180965    0    39907    0      0

 unit 0        1500 192.168.187.0/ 192.168.187.10
fxp1           1514 <Link>      02.00.00.00.00.04  33010497    0 30110800    0      0

 unit 0        1500 10.0.0.0/8     10.0.0.1
                 10.0.0.0/8     10.0.0.4
                 128.0.0.0/2    128.0.0.1
                 128.0.0.0/2    128.0.0.4
               1500 fe80::/64      fe80::200:ff:fe0
                 fec0::/64      fec0::a:0:0:4
               1500                4
lsi            1496 <Link>
dsc             max <Link>                                0    0        0    0      0
lo0             max <Link>                             8980    0     8980    0      0

 unit 0         max 127.0.0.1/8    127.0.0.1
                 192.168.0.1/8  192.168.0.1
 unit 16384     max 127.0.0.1/8    127.0.0.1
 unit 16385     max
gre             max <Link>
ipip            max <Link>
tap             max <Link>
pime            max <Link>
pimd            max <Link>
mtun            max <Link>
so-0/0/0       4474 <Link>                          1679900    0  1068611    0      0

 unit 0        4470 <PtoP>         10.0.60.2              0    0        0    0
 0
so-0/0/1       4474 <Link>                                0    0        0    0      0

 unit 0        4470 <PtoP>         10.0.80.2              0    0        0    0
 0
so-0/0/2       4474 <Link>                                0    0        0    0      0
so-0/0/3       4474 <Link>                                0    0        0    0      0
```


Copyright © 2018, Juniper Networks, Inc.

```
fe-0/1/0    1514 <Link>     00.05.85.88.cc.1f    523120    0   623044    0     0

 unit 0     1500 10.0.90.12/30  10.0.90.14             0    0        0    0
 0
fe-0/1/1    1514 <Link>     00.05.85.88.cc.20        0    0        0    0     0
fe-0/1/2    1514 <Link>     00.05.85.88.cc.21        0    0        0    0     0

...
```

show route forwarding-table interface-name all detail

```
user@host> show route forwarding-table interface-name all detail
Name         Mtu AFam    Network      Address            Ipkts Ierr    Opkts
Oerr  Coll
fxp0        1514         <Link>     00.a0.a5.56.03.83    181005    0    39948
 0     0
 unit 0     1500 inet    192.168.187.0/ 192.168.187.10
fxp1        1514         <Link>     02.00.00.00.00.04  33012676    0 30112468
 0     0
 unit 0     1500 inet    10.0.0.0/8     10.0.0.1
                      10.0.0.0/8     10.0.0.4
                      128.0.0.0/2    128.0.0.1
                      128.0.0.0/2    128.0.0.4
            1500 inet6   fe80::/64      fe80::200:ff:fe0
                      fec0::/64      fec0::a:0:0:4
            1500 tnp                    4
lsi         1496         <Link>
dsc          max         <Link>                              0    0        0
 0     0
lo0          max         <Link>                           8980    0     8980
 0     0
 unit 0      max inet    127.0.0.1/8    127.0.0.1
                      192.168.0.1/8  192.168.0.1
 unit 16384  max inet    127.0.0.1/8    127.0.0.1
 unit 16385  max inet
gre          max         <Link>
ipip         max         <Link>
tap          max         <Link>
pime         max         <Link>
pimd         max         <Link>
mtun         max         <Link>
so-0/0/0    4474         <Link>                        1679980    0  1068661
 0     0
 unit 0     4470 inet    <PtoP>        10.0.60.2               0    0        0
   0     0

...
```

show route hidden

Syntax

show route hidden
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display only hidden route information. A hidden route is unusable, even if it is the best path.

Options

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- *Understanding Hidden Routes*

List of Sample Output

show route hidden on page 2288
show route hidden detail on page 2289
show route hidden extensive on page 2289
show route hidden terse on page 2289

Output Fields

For information about output fields, see the output field table for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route hidden

```
user@host> show route hidden
inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
127.0.0.1/32        [Direct/0] 04:26:38
                    > via lo0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.5.5.5/32         [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
                      AS path: 100 I
```

 Copyright © 2018, Juniper Networks, Inc.

```
                    Unusable
10.12.1.0/24       [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
                    AS path: 100 I
                    Unusable
10.12.80.4/30      [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
                    AS path: I
                    Unusable
...
```

show route hidden detail

```
user@host> show route hidden detail

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
127.0.0.1/32 (1 entry, 0 announced)
         Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Hidden Martian Int>
                Local AS:     1
                Age: 4:27:37
                Task: IF
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete

10.5.5.5/32 (1 entry, 0 announced)
         BGP    Preference: 170/-101
                Route Distinguisher: 10.4.4.4:4
                Next hop type: Unusable
                Next-hop reference count: 6
                State: <Secondary Hidden Int Ext>
                Local AS:     1 Peer AS:     1
                Age: 3:45:09
                Task: BGP_1.10.4.4.4+2493
                AS path: 100 I
                Communities: target:1:999
                VPN Label: 100064
                Localpref: 100
                Router ID: 10.4.4.4
                Primary Routing Table bgp.l3vpn.0

...
```

show route hidden extensive

The output for the **show route hidden extensive** command is identical to that of the **show route hidden detail** command. For sample output, see show route hidden detail on page 2289.

show route hidden terse

```
user@host> show route hidden terse
```

```
inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  127.0.0.1/32       D   0                       >lo0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.5.5.5/32        B 170        100             Unusable        100 I
  10.12.1.0/24       B 170        100             Unusable        100 I
  10.12.80.4/30      B 170        100             Unusable        I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.4.4.4:4:10.5.5.5/32
                     B 170        100             Unusable        100 I
  10.4.4.4:4:10.12.1.0/24
                     B 170        100             Unusable        100 I
  10.4.4.4:4:10.12.80.4/30
                     B 170        100             Unusable        I

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

 Copyright © 2018, Juniper Networks, Inc.

show route inactive-path

List of Syntax

Syntax on page 2291
Syntax (EX Series Switches) on page 2291

Syntax

show route inactive-path
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route inactive-path
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display routes for destinations that have no active route. An inactive route is a route that was not selected as the best path.

Options

none—Display all inactive routes.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- show route active-path on page 2190

List of Sample Output

show route inactive-path on page 2291
show route inactive-path detail on page 2292
show route inactive-path extensive on page 2293
show route inactive-path terse on page 2293

Output Fields

For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route inactive-path

```
user@host> show route inactive-path

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
```

Copyright © 2018, Juniper Networks, Inc.

```
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.100.12/30     [OSPF/10] 03:57:28, metric 1
                    > via so-0/3/0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/8          [Direct/0] 04:39:56
                    > via fxp1.0

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.80.0/30       [BGP/170] 04:38:17, localpref 100
                      AS path: 100 I
                    > to 10.12.80.1 via ge-6/3/2.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route inactive-path detail

```
user@host> show route inactive-path detail

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete

10.12.100.12/30 (2 entries, 1 announced)
         OSPF   Preference: 10
                Next-hop reference count: 1
                Next hop: via so-0/3/0.0, selected
                State: <Int>
                Inactive reason: Route Preference
                Local AS:     1
                Age: 3:58:24    Metric: 1
                Area: 0.0.0.0
                Task: OSPF
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

10.0.0.0/8 (2 entries, 0 announced)
         Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via fxp1.0, selected
                State: <NotBest Int>
```

 Copyright © 2018, Juniper Networks, Inc.

```
                Inactive reason: No difference
                Age: 4:40:52
                Task: IF
                AS path: I

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete

10.12.80.0/30 (2 entries, 1 announced)
         BGP    Preference: 170/-101
                Next-hop reference count: 6
                Source: 10.12.80.1
                Next hop: 10.12.80.1 via ge-6/3/2.0, selected
                State: <Ext>
                Inactive reason: Route Preference
                Peer AS:   100
                Age: 4:39:13
                Task: BGP_100.10.12.80.1+179
                AS path: 100 I
                Localpref: 100
                Router ID: 10.0.0.0
```

show route inactive-path extensive

The output for the **show route inactive-path extensive** command is identical to that of the **show route inactive-path detail** command. For sample output, see show route inactive-path detail on page 2292.

show route inactive-path terse

```
user@host> show route inactive-path terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.12.100.12/30    O  10          1            >so-0/3/0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.0.0.0/8         D   0                       >fxp1.0

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.12.80.0/30      B 170        100            >10.12.80.1      100 I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
```

```
bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```


Copyright © 2018, Juniper Networks, Inc.

show route inactive-prefix

List of Syntax

Syntax on page 2295
Syntax (EX Series Switches) on page 2295

Syntax

show route inactive-prefix
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route inactive-prefix
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display inactive route destinations in each routing table.

Options

none—Display all inactive route destination.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show route inactive-prefix on page 2295
show route inactive-prefix detail on page 2296
show route inactive-prefix extensive on page 2296
show route inactive-prefix terse on page 2296

Output Fields

For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route inactive-prefix

```
user@host> show route inactive-prefix

inet.0: 14 destinations, 14 routes (13 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

127.0.0.1/32        [Direct/0] 00:04:54
                    > via lo0.0
```

Copyright © 2018, Juniper Networks, Inc.

show route inactive-prefix detail

```
user@host> show route inactive-prefix detail

inet.0: 14 destinations, 14 routes (13 active, 0 holddown, 1 hidden)
127.0.0.1/32 (1 entry, 0 announced)
         Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Hidden Martian Int>
                Age: 4:51
                Task: IF
                AS path: I00:04:54
                    > via lo0.0
```

show route inactive-prefix extensive

The output for the **show route inactive-prefix extensive** command is identical to that of the **show route inactive-path detail** command. For sample output, see show route inactive-prefix detail on page 2296.

show route inactive-prefix terse

```
user@host> show route inactive-prefix terse

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  127.0.0.1/32       D   0                       >lo0.0
```


Copyright © 2018, Juniper Networks, Inc.

show route instance

List of Syntax

Syntax on page 2297
Syntax (EX Series Switches and QFX Series) on page 2297

Syntax

show route instance
<brief | detail | summary>
<*instance-name*>
<logical-system (all | *logical-system-name*)>
<operational>

Syntax (EX Series Switches and QFX Series)

show route instance
<brief | detail | summary>
<*instance-name*>
<operational>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display routing instance information.

Options

none—(Same as **brief**) Display standard information about all routing instances.

brief | detail | summary—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**. (These options are not available with the **operational** keyword.)

instance-name—(Optional) Display information for all routing instances whose name begins with this string (for example, **cust1**, **cust11**, and **cust111** are all displayed when you run the **show route instance cust1** command).

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

operational—(Optional) Display operational routing instances.

Required Privilege Level

view

Related Documentation

- *Example: Transporting IPv6 Traffic Across IPv4 Using Filter-Based Tunneling*
- *Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart*

List of Sample Output

show route instance on page 2299
show route instance detail (Graceful Restart Complete) on page 2299
show route instance detail (Graceful Restart Incomplete) on page 2301

show route instance detail (VPLS Routing Instance) on page 2303
show route instance operational on page 2303
show route instance summary on page 2303

Output Fields Table 206 on page 2298 lists the output fields for the **show route instance** command. Output fields are listed in the approximate order in which they appear.

Table 206: show route instance Output Fields

Field Name	Field Description	Level of Output
Instance or *instance-name*	Name of the routing instance.	All levels
Operational Routing Instances	(**operational** keyword only) Names of all operational routing instances.	—
Type	Type of routing instance: **forwarding**, **l2vpn**, **no-forwarding**, **vpls**, **virtual-router**, or **vrf**.	All levels
State	State of the routing instance: **active** or **inactive**.	**brief detail** none
Interfaces	Name of interfaces belonging to this routing instance.	**brief detail** none
Restart State	Status of graceful restart for this instance: **Pending** or **Complete**.	**detail**
Path selection timeout	Maximum amount of time, in seconds, remaining until graceful restart is declared complete. The default is **300**.	**detail**
Tables	Tables (and number of routes) associated with this routing instance.	**brief detail** none
Route-distinguisher	Unique route distinguisher associated with this routing instance.	**detail**
Vrf-import	VPN routing and forwarding instance import policy name.	**detail**
Vrf-export	VPN routing and forwarding instance export policy name.	**detail**
Vrf-import-target	VPN routing and forwarding instance import target community name.	**detail**
Vrf-export-target	VPN routing and forwarding instance export target community name.	**detail**
Vrf-edge-protection-id	Context identifier configured for edge-protection.	**detail**
Fast-reroute-priority	Fast reroute priority setting for a VPLS routing instance: **high**, **medium**, or **low**. The default is **low**.	**detail**
Restart State	Restart state: • **Pending:***protocol-name*—List of protocols that have not yet completed graceful restart for this routing table. • **Complete**—All protocols have restarted for this routing table.	**detail**

 Copyright © 2018, Juniper Networks, Inc.

Table 206: show route instance Output Fields (continued)

Field Name	Field Description	Level of Output
Primary rib	Primary table for this routing instance.	**brief** none **summary**
Active/holddown/hidden	Number of active, hold-down, and hidden routes.	All levels

Sample Output

show route instance

```
user@host> show route instance
Instance             Type
         Primary RIB                                     Active/holddown/hidden
master               forwarding
         inet.0                                          16/0/1
         iso.0                                           1/0/0
         mpls.0                                          0/0/0
         inet6.0                                         2/0/0
         l2circuit.0                                     0/0/0
__juniper_private1__ forwarding
         __juniper_private1__.inet.0                     12/0/0
         __juniper_private1__.inet6.0                    1/0/0
```

show route instance detail (Graceful Restart Complete)

```
user@host> show route instance detail
master:
  Router ID: 10.255.14.176
  Type: forwarding        State: Active
  Restart State: Complete Path selection timeout: 300
  Tables:
    inet.0                 : 17 routes (15 active, 0 holddown, 1 hidden)
    Restart Complete
    inet.3                 : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    iso.0                  : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0                 : 19 routes (19 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l3vpn.0            : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Complete
    inet6.0                : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0            : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
BGP-INET:
  Router ID: 10.69.103.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.103
  Route-distinguisher: 10.255.14.176:103
  Vrf-import: [ BGP-INET-import ]
  Vrf-export: [ BGP-INET-export ]
  Tables:
    BGP-INET.inet.0        : 4 routes (4 active, 0 holddown, 0 hidden)
```

Copyright © 2018, Juniper Networks, Inc.

```
    Restart Complete
BGP-L:
  Router ID: 10.69.104.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.104
  Route-distinguisher: 10.255.14.176:104
  Vrf-import: [ BGP-L-import ]
  Vrf-export: [ BGP-L-export ]
  Tables:
    BGP-L.inet.0           : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Complete
    BGP-L.mpls.0           : 3 routes (3 active, 0 holddown, 0 hidden)
    Restart Complete
L2VPN:
  Router ID: 0.0.0.0
  Type: l2vpn             State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.512
  Route-distinguisher: 10.255.14.176:512
  Vrf-import: [ L2VPN-import ]
  Vrf-export: [ L2VPN-export ]
  Tables:
    L2VPN.l2vpn.0          : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
LDP:
  Router ID: 10.69.105.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.105
  Route-distinguisher: 10.255.14.176:105
  Vrf-import: [ LDP-import ]
  Vrf-export: [ LDP-export ]
  Tables:
    LDP.inet.0             : 5 routes (4 active, 0 holddown, 0 hidden)
    Restart Complete
OSPF:
  Router ID: 10.69.101.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.101
  Route-distinguisher: 10.255.14.176:101
  Vrf-import: [ OSPF-import ]
  Vrf-export: [ OSPF-export ]
  Vrf-import-target: [ target:11111
  Tables:
    OSPF.inet.0            : 8 routes (7 active, 0 holddown, 0 hidden)
    Restart Complete
RIP:
  Router ID: 10.69.102.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.102
  Route-distinguisher: 10.255.14.176:102
  Vrf-import: [ RIP-import ]
  Vrf-export: [ RIP-export ]
```

 Copyright © 2018, Juniper Networks, Inc.

```
  Tables:
    RIP.inet.0            : 6 routes (6 active, 0 holddown, 0 hidden)
    Restart Complete
STATIC:
  Router ID: 10.69.100.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.100
  Route-distinguisher: 10.255.14.176:100
  Vrf-import: [ STATIC-import ]
  Vrf-export: [ STATIC-export ]
  Tables:
    STATIC.inet.0         : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Complete
```

show route instance detail (Graceful Restart Incomplete)

```
user@host> show route instance detail
master:
  Router ID: 10.255.14.176
  Type: forwarding        State: Active
  Restart State: Pending  Path selection timeout: 300
  Tables:
    inet.0                : 17 routes (15 active, 1 holddown, 1 hidden)
    Restart Pending: OSPF LDP
    inet.3                : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: OSPF LDP
    iso.0                 : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0                : 23 routes (23 active, 0 holddown, 0 hidden)
    Restart Pending: LDP VPN
    bgp.l3vpn.0           : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
    inet6.0               : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0           : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
BGP-INET:
  Router ID: 10.69.103.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.103
  Route-distinguisher: 10.255.14.176:103
  Vrf-import: [ BGP-INET-import ]
  Vrf-export: [ BGP-INET-export ]
  Tables:
    BGP-INET.inet.0       : 6 routes (5 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
BGP-L:
  Router ID: 10.69.104.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.104
  Route-distinguisher: 10.255.14.176:104
  Vrf-import: [ BGP-L-import ]
  Vrf-export: [ BGP-L-export ]
  Tables:
    BGP-L.inet.0          : 6 routes (5 active, 0 holddown, 0 hidden)
```

```
    Restart Pending: VPN
    BGP-L.mpls.0          : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
L2VPN:
  Router ID: 0.0.0.0
  Type: l2vpn             State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.512
  Route-distinguisher: 10.255.14.176:512
  Vrf-import: [ L2VPN-import ]
  Vrf-export: [ L2VPN-export ]
  Tables:
    L2VPN.l2vpn.0         : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: VPN L2VPN
LDP:
  Router ID: 10.69.105.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.105
  Route-distinguisher: 10.255.14.176:105
  Vrf-import: [ LDP-import ]
  Vrf-export: [ LDP-export ]
  Tables:
    LDP.inet.0            : 5 routes (4 active, 1 holddown, 0 hidden)
    Restart Pending: OSPF LDP VPN
OSPF:
  Router ID: 10.69.101.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.101
  Route-distinguisher: 10.255.14.176:101
  Vrf-import: [ OSPF-import ]
  Vrf-export: [ OSPF-export ]
  Tables:
    OSPF.inet.0           : 8 routes (7 active, 1 holddown, 0 hidden)
    Restart Pending: OSPF VPN
RIP:
  Router ID: 10.69.102.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.102
  Route-distinguisher: 10.255.14.176:102
  Vrf-import: [ RIP-import ]
  Vrf-export: [ RIP-export ]
  Tables:
    RIP.inet.0            : 8 routes (6 active, 2 holddown, 0 hidden)
    Restart Pending: RIP VPN
STATIC:
  Router ID: 10.69.100.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.100
  Route-distinguisher: 10.255.14.176:100
  Vrf-import: [ STATIC-import ]
  Vrf-export: [ STATIC-export ]
  Tables:
```

 Copyright © 2018, Juniper Networks, Inc.

```
    STATIC.inet.0          : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
```

show route instance detail (VPLS Routing Instance)

```
user@host> show route instance detail test-vpls
test-vpls:
  Router ID: 0.0.0.0
  Type: vpls              State: Active
  Interfaces:
    lsi.1048833
    lsi.1048832
    fe-0/1/0.513
  Route-distinguisher: 10.255.37.65:1
  Vrf-import: [ __vrf-import-test-vpls-internal__ ]
  Vrf-export: [ __vrf-export-test-vpls-internal__ ]
  Vrf-import-target: [ target:300:1 ]
  Vrf-export-target: [ target:300:1 ]
  Vrf-edge-protection-id:  166.1.3.1  Fast-reroute-priority: high
  Tables:
    test-vpls.l2vpn.0          : 3 routes (3 active, 0 holddown, 0 hidden)
```

show route instance operational

```
user@host> show route instance operational
Operational Routing Instances:

master
default
```

show route instance summary

```
user@host> show route instance summary
Instance           Type           Primary rib           Active/holddown/hidden
master             forwarding
                                  inet.0                15/0/1
                                  iso.0                 1/0/0
                                  mpls.0                35/0/0
                                  l3vpn.0               0/0/0
                                  inet6.0               2/0/0
                                  l2vpn.0               0/0/0
                                  l2circuit.0           0/0/0
BGP-INET           vrf
                                  BGP-INET.inet.0       5/0/0
                                  BGP-INET.iso.0        0/0/0
                                  BGP-INET.inet6.0      0/0/0
BGP-L              vrf
                                  BGP-L.inet.0          5/0/0
                                  BGP-L.iso.0           0/0/0
                                  BGP-L.mpls.0          4/0/0
                                  BGP-L.inet6.0         0/0/0
L2VPN              l2vpn
                                  L2VPN.inet.0          0/0/0
                                  L2VPN.iso.0           0/0/0
                                  L2VPN.inet6.0         0/0/0
                                  L2VPN.l2vpn.0         2/0/0
LDP                vrf
                                  LDP.inet.0            4/0/0
                                  LDP.iso.0             0/0/0
                                  LDP.mpls.0            0/0/0
```

Copyright © 2018, Juniper Networks, Inc.

```
                                    LDP.inet6.0          0/0/0
                                    LDP.l2circuit.0      0/0/0
OSPF               vrf
                                    OSPF.inet.0          7/0/0
                                    OSPF.iso.0           0/0/0
                                    OSPF.inet6.0         0/0/0
RIP                vrf
                                    RIP.inet.0           6/0/0
                                    RIP.iso.0            0/0/0
                                    RIP.inet6.0          0/0/0
STATIC             vrf
                                    STATIC.inet.0        4/0/0
                                    STATIC.iso.0         0/0/0
                                    STATIC.inet6.0       0/0/0
```

 Copyright © 2018, Juniper Networks, Inc.

show route label

List of Syntax

Syntax on page 2305
Syntax (EX Series Switches) on page 2305

Syntax

show route label *label*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route label *label*
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.

Description

Display the routes based on a specified Multiprotocol Label Switching (MPLS) label value.

Options

label—Value of the MPLS label.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- *Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs*

List of Sample Output

show route label terse on page 2306
show route label on page 2306
show route label detail on page 2306
show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 2306
show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 2307
show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 2307
show route label detail (Dynamic List Next Hop) on page 2308
show route label extensive on page 2309

Output Fields

For information about output fields, see the output field table for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show route label terse

```
user@host> show route label 100016 terse

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 100016             V 170                       >10.12.80.1
```

show route label

```
user@host> show route label 100016

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
100016             *[VPN/170] 03:25:41
                    > to 10.12.80.1 via ge-6/3/2.0, Pop
```

show route label detail

```
user@host> show route label 100016 detail

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
100016 (1 entry, 1 announced)
        *VPN    Preference: 170
                Next-hop reference count: 2
                Source: 10.12.80.1
                Next hop: 10.12.80.1 via ge-6/3/2.0, selected
                Label operation: Pop
                State: <Active Int Ext>
                Local AS:     1
                Age: 3:23:31
                Task: BGP.0.0.0.0+179
                Announcement bits (1): 0-KRT
                AS path: 100 I
                Ref Cnt: 2
```

show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Next-hop reference count: 3
                Address: 0x9097d90
                Next hop: via vt-0/1/0.1
                Next-hop index: 661
                Label operation: Pop
                Address: 0x9172130
                Next hop: via so-0/0/3.0
                Next-hop index: 654
                Label operation: Swap 299872
```


Copyright © 2018, Juniper Networks, Inc.

```
                State: <Active Int>
                Local AS:  1001
                Age: 8:20       Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
               FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
 src 192.168.142.2
```

show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Next-hop reference count: 3
                Address: 0x9097d90
                Next hop: via vt-0/1/0.1
                Next-hop index: 661
                Label operation: Pop
                Address: 0x9172130
                Next hop: via so-0/0/3.0
                Next-hop index: 654
                Label operation: Swap 299872
                State: **Active Int>
                Local AS:  1001
                Age: 8:20       Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
               FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
 src 192.168.142.2
```

show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show route label 301568 detail

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Address: 0x2735208
                Next-hop reference count: 3
                Next hop type: Router, Next hop index: 1397
                Address: 0x2735d2c
                Next-hop reference count: 3
                Next hop: 1.3.8.2 via ge-1/2/22.0
                Label operation: Pop
                Load balance label:  None;
                Next hop type: Router, Next hop index: 1395
                Address: 0x2736290
                Next-hop reference count: 3
                Next hop: 1.3.4.2 via ge-1/2/18.0
                Label operation: Pop
                Load balance label:  None;
                State: <Active Int AckRequest MulticastRPF>
                Local AS:    10
                Age: 54:05      Metric: 1
                Validation State: unverified
```

```
                    Task: LDP
                    Announcement bits (1): 0-KRT
                    AS path: I
                   FECs bound to route: P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src:
 192.168.219.11
                    Primary Upstream : 1.1.1.3:0--1.1.1.2:0
                       RPF Nexthops  :
                            ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
                            ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
                    Backup Upstream : 1.1.1.3:0--1.1.1.6:0
                       RPF Nexthops :
                            ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xfffe
                            ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xfffe
```

show route label detail (Dynamic List Next Hop)

The output for **show route label detail** shows the two indirect next hop for an ESI.

```
user@host> show route label 299952 detail
mpls.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)
299952 (1 entry, 1 announced)
TSI:
KRT in-kernel 299952 /52 -> {Dyn list:indirect(1048577), indirect(1048574)}
        *EVPN   Preference: 7
                Next hop type: Dynamic List, Next hop index: 1048575
                Address: 0x13f497fc
                Next-hop reference count: 5
                Next hop: ELNH Address 0xb7a3d90 uflags EVPN data
                    Next hop type: Indirect, Next hop index: 0
                    Address: 0xb7a3d90
                    Next-hop reference count: 3
                    Protocol next hop: 10.255.255.2
                    Label operation: Push 301344
                    Indirect next hop: 0x135b5c00 1048577 INH Session ID: 0x181
                        Next hop type: Router, Next hop index: 619
                        Address: 0xb7a3d30
                        Next-hop reference count: 4
                        Next hop: 1.0.0.4 via ge-0/0/1.0
                        Label operation: Push 301344, Push 299792(top)
                        Label TTL action: no-prop-ttl, no-prop-ttl(top)
                        Load balance label: Label 301344: None; Label 299792:
None;
                        Label element ptr: 0xb7a3cc0
                        Label parent element ptr: 0xb7a34e0
                        Label element references: 1
                        Label element child references: 0
                        Label element lsp id: 0
                Next hop: ELNH Address 0xb7a37f0 uflags EVPN data
                    Next hop type: Indirect, Next hop index: 0
                    Address: 0xb7a37f0
                    Next-hop reference count: 3
                    Protocol next hop: 10.255.255.3
                    Label operation: Push 301632
                    Indirect next hop: 0x135b5480 1048574 INH Session ID: 0x180
                        Next hop type: Router, Next hop index: 600
                        Address: 0xb7a3790
                        Next-hop reference count: 4
                        Next hop: 1.0.0.4 via ge-0/0/1.0
                        Label operation: Push 301632, Push 299776(top)
                        Label TTL action: no-prop-ttl, no-prop-ttl(top)
                        Load balance label: Label 301632: None; Label 299776:
```

 Copyright © 2018, Juniper Networks, Inc.

```
None;
                        Label element ptr: 0xb7a3720
                        Label parent element ptr: 0xb7a3420
                        Label element references: 1
                        Label element child references: 0
                        Label element lsp id: 0
                State: <Active Int>
                Age: 1:18
                Validation State: unverified
                Task: evpn global task
                Announcement bits (2): 1-KRT 2-evpn global task
                AS path: I
                Routing Instance blue, Route Type Egress-MAC, ESI
00:11:22:33:44:55:66:77:88:99
```

show route label extensive

The output for the show route label extensive command is identical to that of the **show route label detail** command. For sample output, see show route label detail on page 2306.

show route label-switched-path

List of Syntax
Syntax on page 2310
Syntax (EX Series Switches) on page 2310

Syntax

show route label-switched-path *path-name*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route label-switched-path *path-name*
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.

Description

Display the routes used in an MPLS label-switched path (LSP).

Options

brief | detail | extensive | terse—(Optional) Display the specified level of output.

path-name—LSP tunnel name.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show route label-switched-path on page 2310

Output Fields

For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route label-switched-path

```
user@host> show route label-switched-path sf-to-ny
inet.0: 29 destinations, 29 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32          [MPLS/7] 00:00:06, metric 0
                     > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
3.3.3.3/32         *[MPLS/7] 00:00:06, metric 0
                     > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2.2.2.2/32         *[MPLS/7] 00:00:06, metric 0
                     > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
```

 Copyright © 2018, Juniper Networks, Inc.

```
4.4.4.4/32         *[MPLS/7] 00:00:06, metric 0
                      to 111.222.1.9 via s0-0/0/0, label-switched-path abc
                    > to 111.222.1.9 via s0-0/0/0, label-switched-path xyz
                      to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
111.222.1.9/32      [MPLS/7] 00:00:06, metric 0
                    > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

mpls.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

show route localization

Syntax show route localization

Release Information Command introduced in Junos OS Release 11.4 for T-Series routers.
Command introduced in Junos OS Release 12.3 for MX Series routers.

Description (T320, T640, and T1600 routers only) Display route localization details.

Options **detail**—Display detailed output.

Required Privilege Level view

Related Documentation
- *Example: Configuring Packet Forwarding Engine FIB Localization*

Output Fields Table 207 on page 2312 lists the output fields for the **show route localization** command. Output fields are listed in the approximate order in which they appear.

Table 207: show route localization Output Fields

Field Name	Field Description
FIB-local	FPCs configured as FIB-local.
FIB-remote	FPCs configured as FIB-remote.
Normal	FPCs neither configured as **FIB-local** or **FIB-remote**.
Protocols	IPv4 (inet) or IPv6 (inet6) traffic configured for route localization.

Sample Output

```
user@R0> show route localization
FIB localization ready FPCs (and FIB-local Forwarding Engine addresses)
  FIB-local:   FPC2(4,5)
  FIB-remote:  FPC0, FPC1
  Normal:      FPC3, FPC4, FPC5, FPC6, FPC7
```

```
user@R0> show route localization detail
FIB localization ready FPCs (and FIB-local Forwarding Engine addresses)
  FIB-local:   FPC2(4,5)
  FIB-remote:  FPC0, FPC1
  Normal:      FPC3, FPC4, FPC5, FPC6, FPC7
FIB localization configuration
  Protocols:   inet, inet6
  FIB-local:   FPC2
  FIB-remote:  FPC0, FPC1
```


Copyright © 2018, Juniper Networks, Inc.

```
Forwarding Engine addresses
  FPC0: 1
  FPC1: 2
  FPC2: 4, 5
  FPC3: 6
  FPC4: 8
  FPC5: 11
  FPC6: 13
  FPC7: 15
```

show route martians

List of Syntax

Syntax on page 2314
Syntax (EX Series Switches) on page 2314

Syntax

show route martians
<logical-system (all | *logical-system-name*)>
<table *routing-table-name*>

Syntax (EX Series Switches)

show route martians
<table *routing-table-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display the martian (invalid and ignored) entries associated with each routing table.

Options

none—Display standard information about route martians for all routing tables.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

table *routing-table-name*—(Optional) Display information about route martians for all routing tables whose name begins with this string (for example, **inet.0** and **inet6.0** are both displayed when you run the **show route martians table inet** command).

Required Privilege Level

view

Related Documentation

- *Example: Configuring Class E Martian Addresses for Routing*
- *Understanding Martian Addresses*

List of Sample Output

show route martians on page 2315

Output Fields

Table 208 on page 2314 lists the output fields for the **show route martians** command. Output fields are listed in the approximate order in which they appear

Table 208: show route martians Output Fields

Field Name	Field Description
table-name	Name of the route table in which the route martians reside.
destination-prefix	Route destination.
match value	Route match parameter.

 Copyright © 2018, Juniper Networks, Inc.

Table 208: show route martians Output Fields (continued)

Field Name	Field Description
status	Status of the route: **allowed** or **disallowed**.

Sample Output

show route martians

```
user@host> show route martians

inet.0:
             0.0.0.0/0 exact -- allowed
             0.0.0.0/8 orlonger -- disallowed
             127.0.0.0/8 orlonger -- disallowed
             192.0.0.0/24 orlonger -- disallowed
             240.0.0.0/4 orlonger -- disallowed
             224.0.0.0/4 exact -- disallowed
             224.0.0.0/24 exact -- disallowed

inet.1:
             0.0.0.0/0 exact -- allowed
             0.0.0.0/8 orlonger -- disallowed
             127.0.0.0/8 orlonger -- disallowed
             192.0.0.0/24 orlonger -- disallowed
             240.0.0.0/4 orlonger -- disallowed

inet.2:
             0.0.0.0/0 exact -- allowed
             0.0.0.0/8 orlonger -- disallowed
             127.0.0.0/8 orlonger -- disallowed
             192.0.0.0/24 orlonger -- disallowed
             240.0.0.0/4 orlonger -- disallowed
             224.0.0.0/4 exact -- disallowed
             224.0.0.0/24 exact -- disallowed

inet.3:
             0.0.0.0/0 exact -- allowed
             0.0.0.0/8 orlonger -- disallowed
             127.0.0.0/8 orlonger -- disallowed
             192.0.0.0/24 orlonger -- disallowed
             240.0.0.0/4 orlonger -- disallowed
             224.0.0.0/4 exact -- disallowed
             224.0.0.0/24 exact -- disallowed
...

inet6.0:
             ::1/128 exact -- disallowed
             ff00::/8 exact -- disallowed
             ff02::/16 exact -- disallowed

inet6.1:
             ::1/128 exact -- disallowed

inet6.2:
             ::1/128 exact -- disallowed
             ff00::/8 exact -- disallowed
             ff02::/16 exact -- disallowed
```

Copyright © 2018, Juniper Networks, Inc.

```
inet6.3:
             ::1/128 exact -- disallowed
             ff00::/8 exact -- disallowed
             ff02::/16 exact -- disallowed
...
```

 Copyright © 2018, Juniper Networks, Inc.

show route match-prefix

Syntax show route match-prefix *match-prefix*;

Release Information Command introduced in Junos OS Release 11.4.

Description Allows you to search for routes using regular expressions based on the extended (modern) regular expressions as defined in POSIX 1003.2.

Options ***match-prefix***—Regular expression to match formatted prefix.

Additional Information

Required Privilege Level view

Related Documentation *Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands, Configuration Statements, and Hierarchies*

List of Sample Output
show route match-prefix *:10.255.2.200:6:* (Show all routes matching route distributor 10.255.2.200:6) on page 2317
show route match-prefix 7* (Show all mvpn type-7 routes) on page 2317
show route match-prefix *:224.* (Show all routes matching group 224/4) on page 2317

Output Fields For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route match-prefix *:10.255.2.200:6:* (Show all routes matching route distributor 10.255.2.200:6)

```
user@host> show route match-prefix *:10.255.2.200:6:*
```

show route match-prefix 7* (Show all mvpn type-7 routes)

```
user@host> show route table blue.mvpn.0 match-prefix 7*
Paste
router command output here
```

show route match-prefix *:224.* (Show all routes matching group 224/4)

```
user@host> show route match-prefix *:224.*
```

show route next-hop

List of Syntax

Syntax

show route next-hop *next-hop*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route next-hop *next-hop*
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display the entries in the routing table that are being sent to the specified next-hop address.

Options

brief | detail | extensive | terse—(Optional) Display the specified level of ouput.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

next-hop—Next-hop address.

Required Privilege Level

view

List of Sample Output

Output Fields

For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route next-hop

```
user@host> show route next-hop 192.168.71.254

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16       *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
10.209.0.0/16      *[Static/5] 06:26:25
```

 Copyright © 2018, Juniper Networks, Inc.

```
                       > to 192.168.71.254 via fxp0.0
172.16.0.0/12         *[Static/5] 06:26:25
                       > to 192.168.71.254 via fxp0.0
192.168.0.0/16        *[Static/5] 06:26:25
                       > to 192.168.71.254 via fxp0.0
192.168.102.0/23      *[Static/5] 06:26:25
                       > to 192.168.71.254 via fxp0.0
207.17.136.0/24       *[Static/5] 06:26:25
                       > to 192.168.71.254 via fxp0.0
207.17.136.192/32     *[Static/5] 06:26:25
                       > to 192.168.71.254 via fxp0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route next-hop detail

```
user@host> show route next-hop 192.168.71.254 detail

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
10.10.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

172.16.0.0/12 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
```

```
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

192.168.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

192.168.102.0/23 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

207.17.136.0/24 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

207.17.136.192/32 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
```

 Copyright © 2018, Juniper Networks, Inc.

```
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route next-hop extensive

```
user@host> show route next-hop 192.168.71.254 extensive

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

172.16.0.0/12 (1 entry, 1 announced)
TSI:
KRT in-kernel 172.16.0.0/12 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

192.168.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 192.168.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I
```

```
192.168.102.0/23 (1 entry, 1 announced)
TSI:
KRT in-kernel 192.168.102.0/23 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

207.17.136.192/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.192/32 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

green.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

red.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route next-hop terse

```
user@host> show route next-hop 192.168.71.254 terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
```

 Copyright © 2018, Juniper Networks, Inc.

```
A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.10.0.0/16       S   5                       >192.168.71.254
* 10.209.0.0/16      S   5                       >192.168.71.254
* 172.16.0.0/12      S   5                       >192.168.71.254
* 192.168.0.0/16     S   5                       >192.168.71.254
* 192.168.102.0/23   S   5                       >192.168.71.254
* 207.17.136.0/24    S   5                       >192.168.71.254
* 207.17.136.192/32  S   5                       >192.168.71.254

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route no-community

List of Syntax

Syntax on page 2324
Syntax (EX Series Switches) on page 2324

Syntax

show route no-community
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route no-community
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display the route entries in each routing table that are not associated with any community.

Options

none—(Same as **brief**) Display the route entries in each routing table that are not associated with any community.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show route no-community on page 2324
show route no-community detail on page 2325
show route no-community extensive on page 2325
show route no-community terse on page 2326

Output Fields

For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route no-community

```
user@host> show route no-community
inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16       *[Static/5] 00:36:27
                    > to 192.168.71.254 via fxp0.0
10.209.0.0/16      *[Static/5] 00:36:27
                    > to 192.168.71.254 via fxp0.0
10.255.71.52/32    *[Direct/0] 00:36:27
```

 Copyright © 2018, Juniper Networks, Inc.

```
                      > via lo0.0
10.255.71.63/32      *[OSPF/10] 00:04:39, metric 1
                      > to 35.1.1.2 via ge-3/1/0.0
10.255.71.64/32      *[OSPF/10] 00:00:08, metric 2
                      > to 35.1.1.2 via ge-3/1/0.0
10.255.71.240/32     *[OSPF/10] 00:05:04, metric 2
                        via so-0/1/2.0
                      > via so-0/3/2.0
10.255.71.241/32     *[OSPF/10] 00:05:14, metric 1
                      > via so-0/1/2.0
10.255.71.242/32     *[OSPF/10] 00:05:19, metric 1
                      > via so-0/3/2.0
172.16.12.0/24          *[OSPF/10] 00:05:14, metric 2
                      > via so-0/3/2.0
172.16.14.0/24          *[OSPF/10] 00:00:08, metric 3
                      > to 35.1.1.2 via ge-3/1/0.0
                        via so-0/1/2.0
                        via so-0/3/2.0
172.16.16.0/24          *[OSPF/10] 00:05:14, metric 2
                      > via so-0/1/2.0
.....
```

show route no-community detail

```
user@host> show route no-community detail

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 38:08
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 38:08
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I
....
```

show route no-community extensive

```
user@host> show route no-community extensive

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
```

```
                Local AS:    69
                Age: 2:03:33
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:03:33
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I
```

show route no-community terse

```
user@host> show route no-community terse

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.10.0.0/16       S   5                       >192.168.71.254
* 10.209.0.0/16      S   5                       >192.168.71.254
* 10.255.71.52/32    D   0                       >lo0.0
* 10.255.71.63/32    O  10          1            >35.1.1.2
* 10.255.71.64/32    O  10          2            >35.1.1.2
* 10.255.71.240/32   O  10          2             so-0/1/2.0
                                                 >so-0/3/2.0
* 10.255.71.241/32   O  10          1            >so-0/1/2.0
* 10.255.71.242/32   O  10          1            >so-0/3/2.0
* 172.16.12.0/24        O  10          2            >so-0/3/2.0
* 172.16.14.0/24        O  10          3            >35.1.1.2
                                                  so-0/1/2.0
                                                  so-0/3/2.0
* 172.16.16.0/24        O  10          2            >so-0/1/2.0
...
```


Copyright © 2018, Juniper Networks, Inc.

show route output

List of Syntax

Syntax on page 2327
Syntax (EX Series Switches) on page 2327

Syntax

show route output (address *ip-address* | interface *interface-name*)
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route output (address *ip-address* | interface *interface-name*)
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display the entries in the routing table learned through static routes and interior gateway protocols that are to be sent out the interface with either the specified IP address or specified name.

To view routes advertised to a neighbor or received from a neighbor for the BGP protocol, use the **show route advertising-protocol bgp** and **show route receive-protocol bgp** commands instead.

Options

address ***ip-address***—Display entries in the routing table that are to be sent out the interface with the specified IP address.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

interface ***interface-name***—Display entries in the routing table that are to be sent out the interface with the specified name.

logical-system (all | ***logical-system-name*****)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show route output address on page 2328
show route output address detail on page 2328
show route output address extensive on page 2329
show route output address terse on page 2329
show route output interface on page 2329
show route output interface detail on page 2330
show route output interface extensive on page 2330
show route output interface terse on page 2330

Output Fields For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route output address

```
user@host> show route output address 172.16.36.1/24

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.36.0/24        *[Direct/0] 00:19:56
                    > via so-0/1/2.0
                    [OSPF/10] 00:19:55, metric 1
                    > via so-0/1/2.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route output address detail

```
user@host> show route output address 172.16.36.1 detail

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
172.16.36.0/24 (2 entries, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via so-0/1/2.0, selected
                State: <Active Int>
                Age: 23:00
                Task: IF
                AS path: I
         OSPF   Preference: 10
                Next-hop reference count: 1
                Next hop: via so-0/1/2.0, selected
                State: <Int>
                Inactive reason: Route Preference
                Age: 22:59      Metric: 1
                Area: 0.0.0.0
                Task: OSPF
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```


Copyright © 2018, Juniper Networks, Inc.

```
private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route output address extensive

The output for the **show route output address extensive** command is identical to that of the **show route output address detail** command. For sample output, see show route output address detail on page 2328.

show route output address terse

```
user@host> show route output address 172.16.36.1 terse

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 172.16.36.0/24        D    0                      >so-0/1/2.0
                     O  10           1           >so-0/1/2.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route output interface

```
user@host> show route output interface so-0/1/2.0

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.71.240/32   *[OSPF/10] 00:13:00, metric 2
                      via so-0/1/2.0
                    > via so-0/3/2.0
10.255.71.241/32   *[OSPF/10] 00:13:10, metric 1
                    > via so-0/1/2.0
172.16.14.0/24        *[OSPF/10] 00:05:11, metric 3
                      to 35.1.1.2 via ge-3/1/0.0
                    > via so-0/1/2.0
                      via so-0/3/2.0
172.16.16.0/24        *[OSPF/10] 00:13:10, metric 2
                    > via so-0/1/2.0
172.16.36.0/24        *[Direct/0] 00:13:21
                    > via so-0/1/2.0
                    [OSPF/10] 00:13:20, metric 1
                    > via so-0/1/2.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
```

Copyright © 2018, Juniper Networks, Inc.

```
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route output interface detail

user@host> **show route output interface so-0/1/2.0 detail**

```
inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
10.255.71.240/32 (1 entry, 1 announced)
        *OSPF   Preference: 10
                Next-hop reference count: 2
                Next hop: via so-0/1/2.0
                Next hop: via so-0/3/2.0, selected
                State: <Active Int>
                Age: 14:52      Metric: 2
                Area: 0.0.0.0
                Task: OSPF
                Announcement bits (1): 0-KRT
                AS path: I

10.255.71.241/32 (1 entry, 1 announced)
        *OSPF   Preference: 10
                Next-hop reference count: 4
                Next hop: via so-0/1/2.0, selected
                State: <Active Int>
                Age: 15:02      Metric: 1
                Area: 0.0.0.0
                Task: OSPF
                Announcement bits (1): 0-KRT
                AS path: I
...
```

show route output interface extensive

The output for the **show route output interface extensive** command is identical to that of the **show route output interface detail** command. For sample output, see show route output interface detail on page 2330.

show route output interface terse

user@host> **show route output interface so-0/1/2.0 terse**

```
inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.255.71.240/32   O  10          2              so-0/1/2.0
                                                  >so-0/3/2.0
* 10.255.71.241/32   O  10          1             >so-0/1/2.0
* 172.16.14.0/24        O  10          3              35.1.1.2
                                                  >so-0/1/2.0
                                                   so-0/3/2.0
* 172.16.16.0/24        O  10          2              >so-0/1/2.0
* 172.16.36.0/24        D   0                         >so-0/1/2.0
                     O  10          1             >so-0/1/2.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
```


Copyright © 2018, Juniper Networks, Inc.

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Copyright © 2018, Juniper Networks, Inc.

show route protocol

List of Syntax

Syntax on page 2332
Syntax (EX Series Switches) on page 2332

Syntax

show route protocol *protocol*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route protocol *protocol*
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
ospf2 and **ospf3** options introduced in Junos OS Release 9.2.
ospf2 and **ospf3** options introduced in Junos OS Release 9.2 for EX Series switches.
flow option introduced in Junos OS Release 10.0.
flow option introduced in Junos OS Release 10.0 for EX Series switches.

Description

Display the route entries in the routing table that were learned from a particular protocol.

Options

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

protocol—Protocol from which the route was learned:

- **access**—Access route for use by DHCP application
- **access-internal**—Access-internal route for use by DHCP application
- **aggregate**—Locally generated aggregate route
- **arp**—Route learned through the Address Resolution Protocol
- **atmvpn**—Asynchronous Transfer Mode virtual private network
- **bgp**—Border Gateway Protocol
- **ccc**—Circuit cross-connect
- **direct**—Directly connected route
- **dvmrp**—Distance Vector Multicast Routing Protocol
- **esis**—End System-to-Intermediate System
- **flow**—Locally defined flow-specification route
- **frr**—Precomputed protection route or backup route used when a link goes down
- **isis**—Intermediate System-to-Intermediate System

Copyright © 2018, Juniper Networks, Inc.

- **ldp**—Label Distribution Protocol
- **l2circuit**—Layer 2 circuit
- **l2vpn**—Layer 2 virtual private network
- **local**—Local address
- **mpls**—Multiprotocol Label Switching
- **msdp**—Multicast Source Discovery Protocol
- **ospf**—Open Shortest Path First versions 2 and 3
- **ospf2**—Open Shortest Path First versions 2 only
- **ospf3**—Open Shortest Path First version 3 only
- **pim**—Protocol Independent Multicast
- **rip**—Routing Information Protocol
- **ripng**—Routing Information Protocol next generation
- **rsvp**—Resource Reservation Protocol
- **rtarget**—Local route target virtual private network
- **static**—Statically defined route
- **tunnel**—Dynamic tunnel
- **vpn**—Virtual private network

NOTE: EX Series switches run a subset of these protocols. See the switch CLI for details.

Required Privilege Level view

List of Sample Output

show route protocol l2vpn extensive on page 2340
show route protocol ldp on page 2340
show route protocol ldp extensive on page 2341
show route protocol ospf (Layer 3 VPN) on page 2342
show route protocol ospf detail on page 2343
show route protocol rip on page 2343
show route protocol rip detail on page 2343
show route protocol ripng table inet6 on page 2343
show route protocol static detail on page 2344

Output Fields For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route protocol access

```
user@host> show route protocol access
inet.0: 30380 destinations, 30382 routes (30379 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

13.160.0.3/32      *[Access/13] 00:00:09
                    > to 13.160.0.2 via fe-0/0/0.0
13.160.0.4/32      *[Access/13] 00:00:09
                    > to 13.160.0.2 via fe-0/0/0.0
13.160.0.5/32      *[Access/13] 00:00:09
                    > to 13.160.0.2 via fe-0/0/0.0
```

show route protocol access-internal extensive

```
user@host> show route protocol access-internal 13.160.0.19 extensive
inet.0: 100020 destinations, 100022 routes (100019 active, 0 holddown, 1 hidden)
13.160.0.19/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.19/32 -> {13.160.0.2}
        *Access-internal Preference: 12
                Next-hop reference count: 200000
                Next hop: 13.160.0.2 via fe-0/0/0.0, selected
                State: <Active Int>
       Age: 36
                Task: RPD Unix Domain Server./var/run/rpd_serv.local
                Announcement bits (1): 0-KRT
                AS path: I
```

show route protocol arp

```
user@host> show route protocol arp
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
```

 Copyright © 2018, Juniper Networks, Inc.

```
                      Unusable
20.20.1.4/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.5/32        [ARP/4294967293] 00:04:32, from 20.20.1.1
                      Unusable
20.20.1.6/32        [ARP/4294967293] 00:04:34, from 20.20.1.1
                      Unusable
20.20.1.7/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.8/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.9/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.10/32       [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.11/32       [ARP/4294967293] 00:04:33, from 20.20.1.1
                      Unusable
20.20.1.12/32       [ARP/4294967293] 00:04:33, from 20.20.1.1
                      Unusable
20.20.1.13/32       [ARP/4294967293] 00:04:33, from 20.20.1.1
                      Unusable
...
```

show route protocol bgp

```
user@host> show route protocol bgp 192.168.64.0/21
inet.0: 335832 destinations, 335833 routes (335383 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.64.0/21       *[BGP/170] 6d 10:41:16, localpref 100, from 192.168.69.71
                      AS path: 10458 14203 2914 4788 4788 I
                    > to 192.168.167.254 via fxp0.0
```

show route protocol bgp detail

```
user@host> show route protocol bgp 66.117.63.0/24 detail
inet.0: 335805 destinations, 335806 routes (335356 active, 0 holddown, 450 hidden)
66.117.63.0/24  (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Indirect
                Next-hop reference count: 1006436
                Source: 192.168.69.71
                Next hop type: Router, Next hop index: 324
                Next hop: 192.168.167.254 via fxp0.0, selected
                Protocol next hop: 192.168.69.71
                Indirect next hop: 8e166c0 342
                State: <Active Ext>
                Local AS:    69 Peer AS: 10458
                Age: 6d 10:42:42        Metric2: 0
                Task: BGP_10458.192.168.69.71+179
                Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree
1
                AS path: 10458 14203 2914 4788 4788 I
                Communities: 2914:410 2914:2403 2914:3400
                Accepted
                Localpref: 100
                Router ID: 207.17.136.192
```

show route protocol bgp detail (Labeled Unicast)

```
user@host> show route protocol bgp 1.1.1.8/32 detail
inet.0: 45 destinations, 46 routes (45 active, 0 holddown, 0 hidden)
1.1.1.8/32 (2 entries, 2 announced)
State:
*BGP Preference: 1/-101
Next hop type: Indirect, Next hop index: 0
Address: 0xc007f30
Next-hop reference count: 2
Source: 1.1.1.1
Next hop type: Router, Next hop index: 614
Next hop: 20.1.1.2 via ge-0/0/1.0, selected
Label-switched-path lsp1
Label operation: Push 1000126, Push 1000125, Push 1000124, Push 1000123, Push
299872(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 1000126: None; Label 1000125: None; Label 1000124: None;
 Label 1000123: None; Label 299872: None;
Label element ptr: 0xc007860
Label parent element ptr: 0xc0089a0
Label element references: 1
Label element child references: 0
Label element lsp id: 0
Session Id: 0x140
Protocol next hop: 1.1.1.4
Label operation: Push 1000126, Push 1000125, Push 1000124, Push 1000123(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl, prop-ttl
Load balance label: Label 1000126: None; Label 1000125: None; Label 1000124: None;
 Label 1000123: None;
Indirect next hop: 0xae8d300 1048576 INH Session ID: 0x142
State:
Local AS: 5 Peer AS: 5
Age: 22:43 Metric2: 2
Validation State: unverified
Task: BGP_5.1.1.1.1
Announcement bits (2): 0-KRT 7-Resolve tree 2
AS path: I
Accepted
Route Labels: 1000123(top) 1000124 1000125 1000126
Localpref: 100
Router ID: 1.1.1.1
```

show route protocol bgp detail (Aggregate Extended Community Bandwidth)

```
user@host> show route 10.0.2.0 protocol bgp detail
inet.0: 20 destinations, 26 routes (20 active, 0 holddown, 0 hidden)
10.0.2.0/30 (2 entries, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 0
                Address: 0xb618990
                Next-hop reference count: 3
                Source: 10.0.1.1
                Next hop: 10.0.0.2 via ge-0/0/0.0 balance 40%
                Session Id: 0x0
                Next hop: 10.0.1.1 via ge-0/0/1.0 balance 60%, selected
                Session Id: 0x0
                State: <Active Ext>
                Local AS: 65000 Peer AS: 65001
                Age: 20:33
```


Copyright © 2018, Juniper Networks, Inc.

```
                Validation State: unverified
                Task: BGP_65001.10.0.1.1
                Announcement bits (3): 0-KRT 2-BGP_Listen.0.0.0.0+179
3-BGP_RT_Background
                AS path: 65001 I
                Communities: bandwidth:65000:60000000
          Accepted Multipath
                Localpref: 100
                Router ID: 128.49.121.137
         BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 595
                Address: 0xb7a1330
                Next-hop reference count: 9
                Source: 10.0.0.2
                Next hop: 10.0.0.2 via ge-0/0/0.0, selected
                Session Id: 0x141
                State: <NotBest Ext>
                Inactive reason: Not Best in its group - Active preferred
                Local AS: 65000 Peer AS: 65001
                Age: 20:33
                Validation State: unverified
                Task: BGP_65001.10.0.0.2
                AS path: 65001 I
                Communities: bandwidth:65000:40000000
                Accepted MultipathContrib
                Localpref: 100
                Router ID: 128.49.121.132
```

show route protocol bgp extensive

user@host> **show route protocol bgp 192.168.64.0/21 extensive**

```
inet.0: 335827 destinations, 335828 routes (335378 active, 0 holddown, 450 hidden)
192.168.64.0/21 (1 entry, 1 announced)
TSI:
KRT in-kernel 1.9.0.0/16 -> {indirect(342)}
Page 0 idx 1 Type 1 val db31a80
    Nexthop: Self
    AS path: [69] 10458 14203 2914 4788 4788 I
    Communities: 2914:410 2914:2403 2914:3400
Path 1.9.0.0 from 192.168.69.71 Vector len 4.  Val: 1
        *BGP    Preference: 170/-101
                Next hop type: Indirect
                Next-hop reference count: 1006502
                Source: 192.168.69.71
                Next hop type: Router, Next hop index: 324
                Next hop: 192.168.167.254 via fxp0.0, selected
                Protocol next hop: 192.168.69.71
                Indirect next hop: 8e166c0 342
                State: <Active Ext>
                Local AS:    69 Peer AS: 10458
                Age: 6d 10:44:45        Metric2: 0
                Task: BGP_10458.192.168.69.71+179
                Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree
1
                AS path: 10458 14203 2914 4788 4788 I
                Communities: 2914:410 2914:2403 2914:3400
                Accepted
                Localpref: 100
                Router ID: 207.17.136.192
                Indirect next hops: 1
```

```
                        Protocol next hop: 192.168.69.71
                        Indirect next hop: 8e166c0 342
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 192.168.167.254 via fxp0.0
                        192.168.0.0/16 Originating RIB: inet.0
                          Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: 192.168.167.254 via fxp0.0
```

show route protocol bgp terse

```
user@host> show route protocol bgp 192.168.64.0/21 terse

inet.0: 24 destinations, 32 routes (23 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  192.168.64.0/21    B 170        100            >172.16.100.1    10023 21 I
```

show route protocol direct

```
user@host> show route protocol direct

inet.0: 335843 destinations, 335844 routes (335394 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.8.0/24         *[Direct/0] 17w0d 10:31:49
                    > via fe-1/3/1.0
10.255.165.1/32    *[Direct/0] 25w4d 04:13:18
                    > via lo0.0
172.16.30.0/24      *[Direct/0] 17w0d 23:06:26
                    > via fe-1/3/2.0
192.168.164.0/22   *[Direct/0] 25w4d 04:13:20
                    > via fxp0.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

47.0005.80ff.f800.0000.0108.0001.0102.5516.5001/152
                   *[Direct/0] 25w4d 04:13:21
                    > via lo0.0

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::10:255:165:1/128
                   *[Direct/0] 25w4d 04:13:21
                    > via lo0.0
fe80::2a0:a5ff:fe12:ad7/128
                   *[Direct/0] 25w4d 04:13:21
                    > via lo0.0
```

show route protocol frr

```
user@host> show route protocol frr
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
```

 Copyright © 2018, Juniper Networks, Inc.

```
cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.3 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.4/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.4 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.5/32       *[FRR/200] 00:05:35, from 20.20.1.1
                    > to 20.20.1.5 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.6/32       *[FRR/200] 00:05:37, from 20.20.1.1
                    > to 20.20.1.6 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.7/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.7 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.8/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.8 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.9/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.9 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.10/32      *[FRR/200] 00:05:38, from 20.20.1.1
...
```

show route protocol l2circuit detail

```
user@host> show route protocol l2circuit detail

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
100000 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                Next hop: via ge-2/0/0.0, selected
                Label operation: Pop       Offset: 4
                State: <Active Int>
                Local AS:    99
                Age: 9:52
                Task: Common L2 VC
                Announcement bits (1): 0-KRT
                AS path: I

ge-2/0/0.0 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                Next hop: via so-1/1/2.0 weight 1, selected
                Label-switched-path my-lsp
                Label operation: Push 100000, Push 100000(top)[0] Offset: -4
                Protocol next hop: 10.245.255.63
                Push 100000 Offset: -4
                 Indirect next hop: 86af0c0 298
                State: <Active Int>
                Local AS:    99
                Age: 9:52
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 1-Common L2 VC
                AS path: I

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
```

Copyright © 2018, Juniper Networks, Inc.

```
        *L2CKT  Preference: 7
                Next hop: via so-1/1/2.0 weight 1, selected
                Label-switched-path my-lsp
                Label operation: Push 100000[0]
                Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
                State: <Active Int>
                Local AS:    99
                Age: 10:21
                Task: l2 circuit
                Announcement bits (1): 0-LDP
                AS path: I
                VC Label 100000, MTU 1500, VLAN ID 512
```

show route protocol l2vpn extensive

```
user@host> show route protocol l2vpn extensive

inet.0: 14 destinations, 15 routes (13 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
800001 (1 entry, 1 announced)
TSI:
KRT in-kernel 800001 /36 -> {so-0/0/0.0}
        *L2VPN  Preference: 7
                Next hop: via so-0/0/0.0 weight 49087 balance 97%, selected
                Label operation: Pop      Offset: 4
                State: <Active Int>
                Local AS:    69
                Age: 7:48
                Task: Common L2 VC
                Announcement bits (1): 0-KRT
                AS path: I

so-0/0/0.0 (1 entry, 1 announced)
TSI:
KRT in-kernel so-0/0/0.0.0      /16 -> {indirect(288)}
        *L2VPN  Preference: 7
                Next hop: via so-0/0/1.0, selected
                Label operation: Push 800000 Offset: -4
                Protocol next hop: 10.255.14.220
                Push 800000 Offset: -4
                 Indirect next hop: 85142a0 288
                State: <Active Int>
                Local AS:    69
                Age: 7:48
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 1-Common L2 VC
                AS path: I
                Communities: target:69:1 Layer2-info: encaps:PPP,
                control flags:2, mtu: 0
```

show route protocol ldp

```
user@host> show route protocol ldp
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)
```


Copyright © 2018, Juniper Networks, Inc.

```
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.16.1/32    *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Push 100000
192.168.17.1/32    *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0

private1__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100064             *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Pop
100064(S=0)        *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Pop
100080             *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Swap 100000
```

show route protocol ldp extensive

```
user@host> show route protocol ldp extensive
192.168.16.1/32 (1 entry, 1 announced)
        State: <FlashAll>
        *LDP    Preference: 9
                Next-hop reference count: 3
                Next hop: via t1-4/0/0.0, selected
                Label operation: Push 100000
                State: <Active Int>
                Local AS: 64500
                Age: 1d 23:03:58        Metric: 1
                Task: LDP
                Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
                AS path: I

192.168.17.1/32 (1 entry, 1 announced)
        State: <FlashAll>
        *LDP    Preference: 9
                Next-hop reference count: 3
                Next hop: via t1-4/0/0.0, selected
                State: <Active Int>
                Local AS: 64500
                Age: 1d 23:03:58        Metric: 1
                Task: LDP
                Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
                AS path: I

private1__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

100064 (1 entry, 1 announced)
TSI:
KRT in-kernel 100064 /36 -> {t1-4/0/0.0}
        *LDP    Preference: 9
                Next-hop reference count: 2
                Next hop: via t1-4/0/0.0, selected
                State: <Active Int>
                Local AS: 64500
                Age: 1d 23:03:58        Metric: 1
```

```
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                Prefixes bound to route: 192.168.17.1/32

100064(S=0) (1 entry, 1 announced)
TSI:
KRT in-kernel 100064 /40 -> {t1-4/0/0.0}
        *LDP    Preference: 9
                Next-hop reference count: 2
                Next hop: via t1-4/0/0.0, selected
                Label operation: Pop
                State: <Active Int>
                Local AS: 64500
                Age: 1d 23:03:58        Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I

100080 (1 entry, 1 announced)
TSI:
KRT in-kernel 100080 /36 -> {t1-4/0/0.0}
        *LDP    Preference: 9
                Next-hop reference count: 2
                Next hop: via t1-4/0/0.0, selected
                Label operation: Swap 100000
                State: <Active Int>
                Local AS: 64500
                Age: 1d 23:03:58        Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                Prefixes bound to route: 192.168.16.1/32
```

show route protocol ospf (Layer 3 VPN)

```
user@host> show route protocol ospf
inet.0: 40 destinations, 40 routes (39 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.4/30       *[OSPF/10] 00:05:18, metric 4
                    > via t3-3/2/0.0
10.39.1.8/30        [OSPF/10] 00:05:18, metric 2
                    > via t3-3/2/0.0
10.255.14.171/32   *[OSPF/10] 00:05:18, metric 4
                    > via t3-3/2/0.0
10.255.14.179/32   *[OSPF/10] 00:05:18, metric 2
                    > via t3-3/2/0.0
172.16.233.5/32       *[OSPF/10] 20:25:55, metric 1

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30       [OSPF/10] 00:05:43, metric 1
                    > via so-0/2/2.0
10.255.14.173/32   *[OSPF/10] 00:05:43, metric 1
                    > via so-0/2/2.0
172.16.233.5/32       *[OSPF/10] 20:26:20, metric 1
```

 Copyright © 2018, Juniper Networks, Inc.

show route protocol ospf detail

```
user@host> show route protocol ospf detail
VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30 (2 entries, 0 announced)
         OSPF   Preference:  10
                Nexthop: via so-0/2/2.0, selected
                State: <Int>
                Inactive reason: Route Preference
                Age: 6:25       Metric: 1
                Area: 0.0.0.0
                Task: VPN-AB-OSPF
                AS path: I
                Communities: Route-Type:0.0.0.0:1:0

...
```

show route protocol rip

```
user@host> show route protocol rip
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32   *[RIP/100] 20:24:34, metric 2
                    > to 10.39.1.22 via t3-0/2/2.0
172.16.233.9/32       *[RIP/100] 00:03:59, metric 1
```

show route protocol rip detail

```
user@host> show route protocol rip detail
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32 (1 entry, 1 announced)
        *RIP    Preference: 100
                Nexthop: 10.39.1.22 via t3-0/2/2.0, selected
                State: <Active Int>
                Age: 20:25:02   Metric: 2
                Task: VPN-AB-RIPv2
                Announcement bits (2): 0-KRT 2-BGP.0.0.0.0+179
                AS path: I
                Route learned from 10.39.1.22 expires in 96 seconds
```

show route protocol ripng table inet6

```
user@host> show route protocol ripng table inet6
inet6.0: 4215 destinations, 4215 routes (4214 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1111::1/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::2/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::3/128        *[RIPng/100] 02:13:33, metric 2
```

```
                     > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::4/128         *[RIPng/100] 02:13:33, metric 2
                     > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::5/128         *[RIPng/100] 02:13:33, metric 2
                     > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::6/128         *[RIPng/100] 02:13:33, metric 2
                     > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
```

show route protocol static detail

```
user@host> show route protocol static detail
inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
10.5.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next hop type: Router, Next hop index: 324
                Address: 0x9274010
                Next-hop reference count: 27
                Next hop: 192.168.187.126 via fxp0.0, selected
                Session Id: 0x0
                State: <Active NoReadvrt Int Ext>
                Age: 7w3d 21:24:25
                Validation State: unverified
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.10.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next hop type: Router, Next hop index: 324
                Address: 0x9274010
                Next-hop reference count: 27
                Next hop: 192.168.187.126 via fxp0.0, selected
                Session Id: 0x0
                State: <Active NoReadvrt Int Ext>
                Age: 7w3d 21:24:25
                Validation State: unverified
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.13.10.0/23 (1 entry, 1 announced)
        *Static Preference: 5
                Next hop type: Router, Next hop index: 324
                Address: 0x9274010
                Next-hop reference count: 27
                Next hop: 192.168.187.126 via fxp0.0, selected
                Session Id: 0x0
                State: <Active NoReadvrt Int Ext>
                Age: 7w3d 21:24:25
                Validation State: unverified
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I
```


Copyright © 2018, Juniper Networks, Inc.

show route range

List of Syntax

Syntax on page 2345
Syntax (EX Series Switches) on page 2345

Syntax

show route range
<brief | detail | extensive | terse>
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route range
<brief | detail | extensive | terse>
<*destination-prefix*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display routing table entries using a prefix range.

Options

none—Display standard information about all routing table entries using a prefix range.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

destination-prefix—Destination and prefix mask for the range.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show route range on page 2345
show route range destination-prefix on page 2346
show route range detail on page 2346
show route range extensive on page 2347
show route range terse on page 2348

Output Fields

For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route range

```
user@host> show route range

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
```

```
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16       *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
10.209.0.0/16      *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
10.255.71.14/32    *[Direct/0] 00:30:01
                    > via lo0.0
172.16.0.0/12      *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
192.168.0.0/16     *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
192.168.64.0/21    *[Direct/0] 00:30:01
                    > via fxp0.0
192.168.71.14/32   *[Local/0] 00:30:01
                      Local via fxp0.0
192.168.102.0/23   *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
...
```

show route range destination-prefix

```
user@host> show route range 192.168.0.0/16

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.0/16     *[Static/5] 00:31:14
                    > to 192.168.71.254 via fxp0.0
192.168.64.0/21    *[Direct/0] 00:31:14
                    > via fxp0.0
192.168.71.14/32   *[Local/0] 00:31:14
                      Local via fxp0.0
192.168.102.0/23   *[Static/5] 00:31:14
                    > to 192.168.71.254 via fxp0.0
```

show route range detail

```
user@host> show route range detail

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 30:05
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 30:05
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I
```

 Copyright © 2018, Juniper Networks, Inc.

```
10.255.71.14/32 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Age: 30:05
                Task: IF
                AS path: I

172.16.0.0/12 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 30:05
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

...
```

show route range extensive

```
user@host> show route range extensive

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 30:17
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 30:17
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.255.71.14/32 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Age: 30:17
                Task: IF
                AS path: I
```

```
172.16.0.0/12 (1 entry, 1 announced)
TSI:
KRT in-kernel 172.16.0.0/12 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 30:17
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

...
```

show route range terse

```
user@host> show route range terse

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.10.0.0/16       S   5                       >192.168.71.254
* 10.209.0.0/16      S   5                       >192.168.71.254
* 10.255.71.14/32    D   0                       >lo0.0
* 172.16.0.0/12      S   5                       >192.168.71.254
* 192.168.0.0/16     S   5                       >192.168.71.254
* 192.168.64.0/21    D   0                       >fxp0.0
* 192.168.71.14/32   L   0                        Local
* 192.168.102.0/23   S   5                       >192.168.71.254
* 207.17.136.0/24    S   5                       >192.168.71.254
* 207.17.136.192/32  S   5                       >192.168.71.254

__juniper_private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.0.0.0/8         D   0                       >fxp2.0
                     D   0                       >fxp1.0
* 10.0.0.4/32        L   0                        Local

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  47.0005.80ff.f800.0000.0108.0001.0102.5507.1014/152
*                    D   0                       >lo0.0

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  abcd::10:255:71:14/128
*                    D   0                       >lo0.0
  fe80::280:42ff:fe11:226f/128
*                    D   0                       >lo0.0

__juniper_private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both
```

 Copyright © 2018, Juniper Networks, Inc.

```
A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  fe80::280:42ff:fe11:226f/128
*                    D   0                       >lo0.16385
```

show route receive-protocol

List of Syntax

Syntax on page 2350
Syntax (EX Series Switches) on page 2350

Syntax

show route receive-protocol *protocol neighbor-address*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)

Syntax (EX Series Switches)

show route receive-protocol *protocol neighbor-address*
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display the routing information as it was received through a particular neighbor using a particular dynamic routing protocol.

Options

brief | detail | extensive | terse—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

protocol neighbor-address—Protocol transmitting the route (**bgp**, **dvmrp**, **msdp**, **pim**, **rip**, or **ripng**) and address of the neighboring router from which the route entry was received.

Additional Information

The output displays the selected routes and the attributes with which they were received, but does not show the effects of import policy on the routing attributes.

Required Privilege Level

view

List of Sample Output

show route receive-protocol bgp on page 2353
show route receive-protocol bgp extensive on page 2353
show route receive-protocol bgp table extensive on page 2354
show route receive-protocol bgp logical-system extensive on page 2354
show route receive-protocol bgp detail (Layer 2 VPN) on page 2355
show route receive-protocol bgp extensive (Layer 2 VPN) on page 2355
show route receive-protocol bgp (Layer 3 VPN) on page 2356
show route receive-protocol bgp detail (Layer 3 VPN) on page 2356
show route receive-protocol bgp detail (Long-Lived Graceful Restart) on page 2357
show route receive-protocol bgp detail (Labeled Unicast) on page 2357
show route receive-protocol bgp extensive (Layer 3 VPN) on page 2358
Show route receive protocol (Segment Routing Traffic Engineering) on page 2358

Copyright © 2018, Juniper Networks, Inc.

Output Fields Table 209 on page 2351 describes the output fields for the **show route receive-protocol** command. Output fields are listed in the approximate order in which they appear.

Table 209: show route receive-protocol Output Fields

Field Name	Field Description	Level of Output
routing-table-name	Name of the routing table—for example, inet.0.	All levels
number destinations	Number of destinations for which there are routes in the routing table.	All levels
number routes	Number of routes in the routing table and total number of routes in the following states: • **active** • **holddown** (routes that are in pending state before being declared inactive) • **hidden** (routes that are not used because of a routing policy)	All levels
Prefix	Destination prefix.	none **brief**
MED	Multiple exit discriminator value included in the route.	none **brief**
destination-prefix (entry, announced)	Destination prefix. The **entry** value is the number of routes for this destination, and the **announced** value is the number of routes being announced for this destination.	**detail extensive**
Accepted LongLivedStale	The LongLivedStale flag indicates that the route was marked LLGR-stale by this router, as part of the operation of LLGR receiver mode. Either this flag or the LongLivedStaleImport flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag.	**detail extensive**
Accepted LongLivedStaleImport	The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy. Either this flag or the LongLivedStale flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag. Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and import into the inet.0 routing table	**detail extensive**
ImportAccepted LongLivedStaleImport	Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and imported into the inet.0 routing table The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy.	**detail extensive**
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.	**detail extensive**
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.	**detail extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 209: show route receive-protocol Output Fields (continued)

Field Name	Field Description	Level of Output
VPN Label	Virtual private network (VPN) label. Packets are sent between CE and PE routing devices by advertising VPN labels. VPN labels transit over either an RSVP or an LDP label-switched path (LSP) tunnel.	**detail extensive**
Next hop	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.	All levels
Localpref or Lclpref	Local preference value included in the route.	All levels
AS path	Autonomous system (AS) path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used the AS-path merge process, as defined in RFC 4893. • **[ ]**—If more than one AS number is configured on the router, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.	All levels
Route Labels	Stack of labels carried in the BGP route update.	**detail extensive**
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.	**detail extensive**
Originator ID	(For route reflected output only) Address of routing device that originally sent the route to the route reflector.	**detail extensive**
Communities	Community path attribute for the route. See the Output Field table in the show route detail command for all possible values for this field.	**detail extensive**
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.	**detail extensive**
Attrset AS	Number, local preference, and path of the AS that originated the route. These values are stored in the **Attrset** attribute at the originating routing device.	**detail extensive**

Copyright © 2018, Juniper Networks, Inc.

Table 209: show route receive-protocol Output Fields (continued)

Field Name	Field Description	Level of Output
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).	**detail extensive**
control flags	Control flags: **none** or **Site Down**.	**detail extensive**
mtu	Maximum transmission unit (MTU) of the Layer 2 circuit.	**detail extensive**

Sample Output

show route receive-protocol bgp

```
user@host> show route receive-protocol bgp 10.255.245.215

inet.0: 28 destinations, 33 routes (27 active, 0 holddown, 1 hidden)
Prefix                   Next hop             MED     Lclpref    AS path
10.22.1.0/24             10.255.245.215       0       100        I
10.22.2.0/24             10.255.245.215       0       100        I
```

show route receive-protocol bgp extensive

```
user@host> show route receive-protocol bgp 10.255.245.63 extensive
inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
Prefix             Next hop                 MED     Lclpref AS path
172.16.1.0/24 (1 entry, 1 announced)
     Next hop: 10.0.50.3
     Localpref: 100
     AS path: I <Orginator>
     Cluster list:  10.2.3.1
     Originator ID: 10.255.245.45
172.16.163.0/16 (1 entry, 1 announced)
     Next hop: 111.222.5.254
     Localpref: 100
     AS path: I <Originator>
     Cluster list:  10.2.3.1
     Originator ID: 10.255.245.68
172.16.164.0/16 (1 entry, 1 announced)
     Next hop: 111.222.5.254
     Localpref: 100
     AS path: I <Originator>
     Cluster list:  10.2.3.1
     Originator ID: 10.255.245.45
172.16.195.0/24 (1 entry, 1 announced)
     Next hop: 111.222.5.254
     Localpref: 100
     AS path: I <Originator>
     Cluster list:  10.2.3.1
     Originator ID: 10.255.245.68
inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
Prefix             Next hop                 MED     Lclpref AS path
inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix             Next hop                 MED     Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Next hop                 MED     Lclpref AS path
mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)
```

Copyright © 2018, Juniper Networks, Inc.

show route receive-protocol bgp table extensive

```
user@host> show route receive-protocol bgp 207.17.136.192 table inet.0 66.117.68.0/24 extensive
inet.0: 227315 destinations, 227316 routes (227302 active, 0 holddown, 13 hidden)
* 66.117.63.0/24 (1 entry, 1 announced)
     Nexthop: 207.17.136.29
     Localpref: 100
     AS path: AS2 PA[6]: 14203 2914 3356 29748 33437 AS_TRANS
     AS path: AS4 PA[2]: 33437 393219
     AS path: Merged[6]: 14203 2914 3356 29748 33437 393219 I
     Communities: 2914:420
```

show route receive-protocol bgp logical-system extensive

```
user@host> show route receive-protocol bgp 10.0.0.9 logical-system PE4 extensive
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)
* 10.0.0.0/30 (1 entry, 1 announced)
     Accepted
     Route Label: 3
     Nexthop: 10.0.0.9
     AS path: 13979 I

* 10.0.0.4/30 (1 entry, 1 announced)
     Accepted
     Route Label: 3
     Nexthop: 10.0.0.9
     AS path: 13979 I

  10.0.0.8/30 (2 entries, 1 announced)
     Accepted
     Route Label: 3
     Nexthop: 10.0.0.9
     AS path: 13979 I

* 10.9.9.1/32 (1 entry, 1 announced)
     Accepted
     Route Label: 3
     Nexthop: 10.0.0.9
     AS path: 13979 I

* 10.100.1.1/32 (1 entry, 1 announced)
     Accepted
     Route Label: 3
     Nexthop: 10.0.0.9
     AS path: 13979 I

* 172.16.44.0/24 (1 entry, 1 announced)
     Accepted
     Route Label: 300096
     Nexthop: 10.0.0.9
     AS path: 13979 I
     AIGP: 203

* 172.16.55.0/24 (1 entry, 1 announced)
     Accepted
     Route Label: 300112
     Nexthop: 10.0.0.9
     AS path: 13979 7018 I
     AIGP: 25
```

 Copyright © 2018, Juniper Networks, Inc.

```
* 172.16.66.0/24 (1 entry, 1 announced)
     Accepted
     Route Label: 300144
     Nexthop: 10.0.0.9
     AS path: 13979 7018 I

* 172.16.99.0/24 (1 entry, 1 announced)
     Accepted
     Route Label: 300160
     Nexthop: 10.0.0.9
     AS path: 13979 7018 I
```

show route receive-protocol bgp detail (Layer 2 VPN)

```
user@host> show route receive-protocol bgp 10.255.14.171 detail
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix             Nexthop                MED    Lclpref AS path
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix             Nexthop                MED    Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Nexthop                MED    Lclpref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix             Nexthop                MED    Lclpref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0
hidden)
Prefix             Nexthop                MED    Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
     Route Distinguisher: 10.255.245.35:1
     Label-base : 800000, range : 4, status-vector : 0x0
     Nexthop: 10.255.245.35
     Localpref: 100
     AS path: I
     Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
     control flags: 0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Nexthop                MED    Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
     Route Distinguisher: 10.255.245.35:1
     Label-base : 800000, range : 4, status-vector : 0x0
     Nexthop: 10.255.245.35
     Localpref: 100
     AS path: I
     Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
     control flags:0, mtu: 0
```

show route receive-protocol bgp extensive (Layer 2 VPN)

```
user@host> show route receive-protocol bgp 10.255.14.171 extensive
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix             Nexthop                MED    Lclpref AS path
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix             Nexthop                MED    Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Nexthop                MED    Lclpref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix             Nexthop                MED    Lclpref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Prefix             Nexthop                MED    Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
     Route Distinguisher: 10.255.245.35:1
```

Copyright © 2018, Juniper Networks, Inc.

```
     Label-base : 800000, range : 4, status-vector : 0x0
     Nexthop: 10.255.245.35
     Localpref: 100
     AS path: I
     Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
     control flags:0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix                Nexthop               MED     Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
     Route Distinguisher: 10.255.245.35:1
     Label-base : 800000, range : 4, status-vector : 0x0
     Nexthop: 10.255.245.35
     Localpref: 100
     AS path: I
     Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
     control flags:0, mtu: 0
```

show route receive-protocol bgp (Layer 3 VPN)

```
user@host> show route receive-protocol bgp 10.255.14.171
      inet.0: 33 destinations, 33 routes (32 active, 0 holddown, 1 hidden)
      Prefix                Nexthop               MED     Lclpref AS path
      inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
      Prefix                Nexthop               MED     Lclpref AS path
      VPN-A.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
      Prefix                Nexthop               MED     Lclpref AS path
      10.255.14.175/32   10.255.14.171                        100 2 I
      10.255.14.179/32   10.255.14.171              2         100 I
      VPN-B.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
      Prefix                Nexthop               MED     Lclpref AS path
      10.255.14.175/32   10.255.14.171                        100 2 I
      10.255.14.177/32   10.255.14.171                        100 I
      iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
      Prefix                Nexthop               MED     Lclpref AS path
      mpls.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
      Prefix                Nexthop               MED     Lclpref AS path
      bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
      Prefix                Nexthop               MED     Lclpref AS path
      10.255.14.171:300:10.255.14.177/32
                         10.255.14.171                        100 I
      10.255.14.171:100:10.255.14.179/32
                         10.255.14.171              2         100 I
      10.255.14.171:200:10.255.14.175/32
                         10.255.14.171                        100 2 I
```

show route receive-protocol bgp detail (Layer 3 VPN)

```
user@host> show route receive-protocol bgp 10.255.14.174 detail
inet.0: 16 destinations, 17 routes (15 active, 0 holddown, 1 hidden)
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
vpna.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
* 10.49.0.0/30 (1 entry, 1 announced)
     Route Distinguisher: 10.255.14.176:2
     VPN Label: 101264
     Nexthop: 10.255.14.174
     Localpref: 100
     AS path: I
     Communities: target:200:100
     AttrSet AS: 100
         Localpref: 100
```


Copyright © 2018, Juniper Networks, Inc.

```
        AS path: I
* 10.255.14.172/32 (1 entry, 1 announced)
     Route Distinguisher: 10.255.14.176:2
     VPN Label: 101280
     Nexthop: 10.255.14.174
     Localpref: 100
     AS path: I
     Communities: target:200:100
     AttrSet AS: 100
         Localpref: 100
         AS path: I
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
* 10.255.14.174:2:10.49.0.0/30 (1 entry, 0 announced)
     Route Distinguisher: 10.255.14.174:2
     VPN Label: 101264
     Nexthop: 10.255.14.174
     Localpref: 100
     AS path: I
     Communities: target:200:100
     AttrSet AS: 100
         Localpref: 100
         AS path: I
* 10.255.14.174:2:10.255.14.172/32 (1 entry, 0 announced)
     Route Distinguisher: 10.255.14.174:2
     VPN Label: 101280
     Nexthop: 10.255.14.174
     Localpref: 100
     AS path: I
     Communities: target:200:100
     AttrSet AS: 100
         Localpref: 100
         AS path: I
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```

show route receive-protocol bgp detail (Long-Lived Graceful Restart)

```
user@host> show route receive-protocol bgp 10.4.12.11 detail

  bgp.l2vpn.0: 38 destinations, 39 routes (37 active, 0 holddown, 1 hidden)
  * 172.16.1.4:100:172.16.1.4/96 AD (1 entry, 1 announced)
    Accepted LongLivedStale LongLivedStaleImport
    Nexthop: 10.4.12.11
    Localpref: 100
    AS path: I
```

show route receive-protocol bgp detail (Labeled Unicast)

```
user@host> show route receive-protocol bgp 1.1.1.1 detail
inet.0: 45 destinations, 46 routes (45 active, 0 holddown, 0 hidden)
* 1.1.1.8/32 (2 entries, 2 announced)
Accepted
Route Labels: 1000123(top) 1000124 1000125 1000126
Nexthop: 1.1.1.4
Localpref: 100
AS path: I
Entropy label capable, next hop field matches route next hop

inet.3: 15 destinations, 21 routes (6 active, 0 holddown, 14 hidden)
```

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)

inet6.0: 26 destinations, 28 routes (26 active, 0 holddown, 0 hidden)

* 100::1/128 (2 entries, 2 announced)
Accepted
Route Labels: 1000123(top) 1000124 1000125 1000126
Nexthop: ::ffff:1.1.1.4
Localpref: 100
AS path: I

inet6.3: 22 destinations, 23 routes (22 active, 0 holddown, 0 hidden)
```

show route receive-protocol bgp extensive (Layer 3 VPN)

```
user@host> show route receive-protocol bgp 10.255.245.63 extensive
inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
     Prefix            Nexthop              MED    Lclpref AS path
     172.16.1.0/24 (1 entry, 1 announced)
          Nexthop: 10.0.50.3
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.45
     172.16.163.0/16 (1 entry, 1 announced)
          Nexthop: 111.222.5.254
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.68
     172.16.164.0/16 (1 entry, 1 announced)
          Nexthop: 111.222.5.254
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.45
     172.16.195.0/24 (1 entry, 1 announced)
          Nexthop: 111.222.5.254
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.68
     inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
     Prefix            Nexthop              MED    Lclpref AS path
     inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
     Prefix            Nexthop              MED    Lclpref AS path
     iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
     Prefix            Nexthop              MED    Lclpref AS path
     mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)
```

Show route receive protocol (Segment Routing Traffic Engineering)

```
show route receive protocol bgp 10.1.1.4
bgp.inetcolor.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

* 50-4.4.4.4-1234<sr6>/96 (1 entry, 0 announced)
     Import Accepted
```

 Copyright © 2018, Juniper Networks, Inc.

```
     Distinguisher: 50
     Color: 1234
     Nexthop: 10.1.1.4
     Localpref: 100
     AS path: 3 I
     Communities: target:1.1.1.1:1


inetcolor.0: 6 destinations, 7 routes (6 active, 0 holddown, 0 hidden)
* 4.4.4.4-1234<c6>/64 (1 entry, 1 announced)
     Import Accepted
     Color: 1234
     Nexthop: 10.1.1.4
     Localpref: 100
     AS path: 3 I
     Communities: target:1.1.1.1:1

user@host# run show route receive-protocol bgp 5001:1::4
bgp.inet6color.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

* 50-2001:1::4-1234<sr6>/192 (1 entry, 0 announced)
     Import Accepted
     Distinguisher: 50
     Color: 1234
     Nexthop: ::ffff:1.1.1.4
     Localpref: 100
     AS path: 3 I
     Communities: target:1.1.1.1:1


inet6color.0: 6 destinations, 7 routes (6 active, 0 holddown, 0 hidden)
* 2001::5-1234<c6>/160 (1 entry, 1 announced)
     Import Accepted
     Color: 1234
     Nexthop: ::ffff:1.1.1.5
     Localpref: 100
     AS path: 3 I
     Communities: target:2:1
```

show route resolution

List of Syntax

Syntax on page 2360
Syntax (EX Series Switches) on page 2360

Syntax

```
show route resolution
<brief | detail | extensive | summary>
<index index>
<logical-system (all | logical-system-name)>
<prefix>
<table routing-table-name>
<unresolved>
```

Syntax (EX Series Switches)

```
show route resolution
<brief | detail | extensive | summary>
<index index>
<prefix>
<table routing-table-name>
<unresolved>
```

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display the entries in the next-hop resolution database. This database provides for recursive resolution of next hops through other prefixes in the routing table.

Options

none—Display standard information about all entries in the next-hop resolution database.

brief | detail | extensive | summary—(Optional) Display the specified level of output.

index *index*—(Optional) Show the index of the resolution tree.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

prefix* *network/destination-prefix—(Optional) Display database entries for the specified address.

table *routing-table-name*—(Optional) Display information about a particular routing table (for example, **inet.0**) where policy-based export is currently enabled.

unresolved—(Optional) Display routes that could not be resolved.

Required Privilege Level

view

Related Documentation

- *Example: Configuring Route Resolution on PE Routers*

Copyright © 2018, Juniper Networks, Inc.

List of Sample Output

Output Fields Table 210 on page 2362 describes the output fields for the **show route resolution** command. Output fields are listed in the approximate order in which they appear.

Table 210: show route resolution Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table whose prefixes are resolved using the entries in the route resolution database. For routing table groups, this is the name of the primary routing table whose prefixes are resolved using the entries in the route resolution database.
Tree index	Tree index identifier.
Nodes	Number of nodes in the tree.
Reference count	Number of references made to the next hop.
Contributing routing tables	Routing tables used for next-hop resolution.
Originating RIB	Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of **inet.0** resolving through **inet.0** and **inet.3**, this field indicates which routing table, **inet.0** or **inet.3**, provided the best path for a particular prefix.
Metric	Metric associated with the forwarding next hop.
Node path count	Number of nodes in the path.
Forwarding next hops	Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it. Merged—Merged next hops when recursive resolution of multipath is configured.

Sample Output

show route resolution detail

```
user@host> show route resolution detail
Tree Index: 1, Nodes 0, Reference Count 1
Contributing routing tables: inet.3
Tree Index: 2, Nodes 23, Reference Count 1
Contributing routing tables: inet.0 inet.3
10.10.0.0/16 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.0/30 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.1/32 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 0
10.31.1.4/30 Originating RIB: inet.0
  Node path count: 1
```

 Copyright © 2018, Juniper Networks, Inc.

```
  Forwarding nexthops: 1
10.31.1.5/32 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 0
10.31.2.0/30 Originating RIB: inet.0
  Metric: 2 Node path count: 1
  Forwarding nexthops: 2
10.31.11.0/24 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
```

show route resolution (Multipath Resolution)

```
user@host> show route resolution detail
user@host> show route resolution detail 10.1.1.2
Tree Index: 1, Nodes 36, Reference Count 3
Contributing routing tables: inet.0 inet.3
Policy: [ abc ]
10.1.1.2/32 Originating RIB: inet.0
  Node path count: 1
  Next hop subtype: INDIRECT
  Indirect next hops: 2
        Protocol next hop: 10.1.1.1
        Inode flags: 0x206 path flags: 0x08
        Path fnh link: 0xc9321c0 path inh link: 0x0
        Indirect next hop: 0xb2b20f0 1048574 INH Session ID: 0x143
        Indirect path forwarding next hops: 1
                Next hop type: Router
                Next hop: 12.1.1.2 via ge-2/0/1.0
                Session Id: 0x144
                Next hop: 13.1.1.2 via ge-2/0/2.0
                Session Id: 0x145

 10.1.1.1/32 Originating RIB: inet.0
   Node path count: 1
  Node flags: 1
   Forwarding nexthops: 1 (Merged)
  Nexthop: 12.1.1.2 via ge-2/0/1.0

  Nexthop: 13.1.1.2 via ge-2/0/2.0
```

```
user@host> show route resolution summary
Tree Index: 1, Nodes 7, Reference Count 2
Contributing routing tables: inet.3
Tree Index: 2, Nodes 7, Reference Count 8213
Contributing routing tables: inet.3
Policy: [ RRwM ]
Tree Index: 3, Nodes 7, Reference Count 2
Contributing routing tables: inet6.3
Tree Index: 4, Nodes 1, Reference Count 1
Contributing routing tables: iso.0
Tree Index: 5, Nodes 1000061, Reference Count 13
Contributing routing tables: inet.0 inet.3
Policy: [ Community-RRwM ]
Tree Index: 6, Nodes 2013, Reference Count 6
Contributing routing tables: inet6.0 inet6.3
Policy: [ RRwM ]
Tree Index: 7, Nodes 7, Reference Count 1501
Contributing routing tables: inet6.3
Policy: [ RRwM ]
```

```
Tree Index: 8, Nodes 1000061, Reference Count 2
Contributing routing tables: inet.0 inet.3
Policy: [ RRwM ]
```

show route resolution summary

```
user@host> show route resolution summary
Tree Index: 1, Nodes 24, Reference Count 1
 Contributing routing tables: :voice.inet.0 :voice.inet.3
 Tree Index: 2, Nodes 2, Reference Count 1
 Contributing routing tables: inet.3
 Tree Index: 3, Nodes 43, Reference Count 1
 Contributing routing tables: inet.0 inet.3
```

show route resolution unresolved

```
user@host> show route resolution unresolved
Tree Index 1
vt-3/2/0.32769.0      /16
        Protocol Nexthop: 10.255.71.238 Push 800000
        Indirect nexthop: 0 -
vt-3/2/0.32772.0      /16
        Protocol Nexthop: 10.255.70.103 Push 800008
        Indirect nexthop: 0 -
Tree Index 2
```


Copyright © 2018, Juniper Networks, Inc.

show route snooping

Syntax

show route snooping
<brief | detail | extensive | terse>
<all>
<best *address/prefix*>
<exact *address*>
<logical-system *logical-system-name*>
<range *prefix-range*>
<summary>
<table *table-name*>

Release Information

Command introduced in Junos OS Release 8.5.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display the entries in the routing table that were learned from snooping.

Options

none—Display the entries in the routing table that were learned from snooping.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

all—(Optional) Display all entries, including hidden entries.

best *address/prefix*—(Optional) Display the longest match for the provided address and optional prefix.

exact *address/prefix*—(Optional) Display exact matches for the provided address and optional prefix.

logical-system *logical-system-name*—(Optional) Display information about a particular logical system, or type ’all’.

range *prefix-range*—(Optional) Display information for the provided address range.

summary—(Optional) Display route snooping summary statisitics.

table *table-name*—(Optional) Display information for the named table.

Required Privilege Level

view

List of Sample Output

show route snooping detail on page 2366
show route snooping logical-system all on page 2366

Output Fields

For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show route snooping detail

```
user@host> show route snooping detail
__+domainAll__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

224.0.0.2/32 (1 entry, 1 announced)
        *IGMP   Preference: 0
                Next hop type: MultiRecv
                Next-hop reference count: 4
                State: <Active NoReadvrt Int>
                Age: 2:24
                Task: IGMP
                Announcement bits (1): 0-KRT
                AS path: I

224.0.0.22/32 (1 entry, 1 announced)
        *IGMP   Preference: 0
                Next hop type: MultiRecv
                Next-hop reference count: 4
                State: <Active NoReadvrt Int>
                Age: 2:24
                Task: IGMP
                Announcement bits (1): 0-KRT
                AS path: I

__+domainAll__.inet.1: 36 destinations, 36 routes (36 active, 0 holddown, 0 hidden)

224.0.0.0.0.0.0.0/24 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4), Next hop index: 1048584
                Next-hop reference count: 4
                State: <Active Int>
                Age: 2:24
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

<snip>
```

show route snooping logical-system all

```
user@host> show route snooping logical-system all

logical-system: default

inet.1: 20 destinations, 20 routes (20 active, 0 holddown, 0 hidden)
Restart Unsupported
+ = Active Route, - = Last Active, * = Both

0.0,0.1,0.0,232.1.1.65,100.1.1.2/112*[Multicast/180] 00:07:36
                      Multicast (IPv4) Composite
0.0,0.1,0.0,232.1.1.66,100.1.1.2/112*[Multicast/180] 00:07:36
                      Multicast (IPv4) Composite
0.0,0.1,0.0,232.1.1.67,100.1.1.2/112*[Multicast/180] 00:07:36

<snip>

default-switch.inet.1: 237 dest, 237 rts (237 active, 0 holddown, 0 hidden)
```

 Copyright © 2018, Juniper Networks, Inc.

```
Restart Complete
+ = Active Route, - = Last Active, * = Both

0.15,0.1,0.0,0.0.0.0,0.0.0.0,2/120*[Multicast/180] 00:08:21
                      Multicast (IPv4) Composite
0.15,0.1,0.0,0.0.0.0,0.0.0.0,2,17/128*[Multicast/180] 00:08:21
                      Multicast (IPv4) Composite

<snip>
```

show route source-gateway

List of Syntax

Syntax on page 2368
Syntax (EX Series Switches) on page 2368

Syntax

show route source-gateway *address*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route source-gateway *address*
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display the entries in the routing table that were learned from a particular address. The **Source** field in the **show route detail** command output lists the source for each route, if known.

Options

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

address—IP address of the system.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show route source-gateway on page 2368
show route source-gateway detail on page 2369
show route source-gateway extensive on page 2371

Output Fields

For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route source-gateway

```
user@host> show route source-gateway 10.255.70.103
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
```

 Copyright © 2018, Juniper Networks, Inc.

```
private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:1:3:1/96
                   *[BGP/170] 12:12:24, localpref 100, from 10.255.70.103
                      AS path: I
                    > via so-0/3/0.0, label-switched-path green-r1-r3

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:2:3:1/96
                   *[BGP/170] 12:12:24, localpref 0, from 10.255.70.103
                      AS path: I
                    > via so-0/3/0.0, label-switched-path green-r1-r3

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:1:3:1/96
                   *[BGP/170] 12:12:24, localpref 100, from 10.255.70.103
                      AS path: I
                    > via so-0/3/0.0, label-switched-path green-r1-r3

10.255.70.103:2:3:1/96
                   *[BGP/170] 12:12:24, localpref 0, from 10.255.70.103
                      AS path: I
                    > via so-0/3/0.0, label-switched-path green-r1-r3
```

show route source-gateway detail

```
user@host> show route source-gateway 10.255.70.103 detail
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
```

```
inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete
green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

Restart Complete
10.255.70.103:1:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.70.103:1
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:14:00   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                Announcement bits (1): 0-green-l2vpn
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Label-base: 800008, range: 8
                Localpref: 100
                Router ID: 10.255.70.103
                Primary Routing Table bgp.l2vpn.0

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:2:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-1
                Route Distinguisher: 10.255.70.103:2
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:14:00   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                Announcement bits (1): 0-red-l2vpn
                AS path: I
                Communities: target:11111:2 Layer2-info: encaps:VPLS,
                control flags:Site-Down, mtu: 0
                Label-base: 800016, range: 8
                Localpref: 0
                Router ID: 10.255.70.103
                Primary Routing Table bgp.l2vpn.0

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.70.103:1
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Active Int Ext>
                Local AS:    69 Peer AS:    69
```

 Copyright © 2018, Juniper Networks, Inc.

```
                Age: 12:14:00   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS, control
flags:, mtu: 0
                Label-base: 800008, range: 8
                Localpref: 100
                Router ID: 10.255.70.103
                Secondary Tables: green.l2vpn.0
10.255.70.103:2:3:1/96 (1 entry, 0 announced)
        *BGP    Preference: 170/-1
                Route Distinguisher: 10.255.70.103:2
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:14:00   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                AS path: I
                Communities: target:11111:2 Layer2-info: encaps:VPLS,
                control flags:Site-Down,
                mtu: 0
                Label-base: 800016, range: 8
                Localpref: 0
                Router ID: 10.255.70.103
                Secondary Tables: red.l2vpn.0
```

show route source-gateway extensive

```
user@host> show route source-gateway 10.255.70.103 extensive
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
10.255.70.103:1:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.70.103:1
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:15:24   Metric2: 1
```

```
                Task: BGP_69.10.255.70.103+179
                Announcement bits (1): 0-green-l2vpn
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Label-base: 800008, range: 8
                Localpref: 100
                Router ID: 10.255.70.103
                Primary Routing Table bgp.l2vpn.0

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:2:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-1
                Route Distinguisher: 10.255.70.103:2
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:15:24   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                Announcement bits (1): 0-red-l2vpn
                AS path: I
                Communities: target:11111:2 Layer2-info: encaps:VPLS,
                control flags:Site-Down, mtu: 0
                Label-base: 800016, range: 8
                Localpref: 0
                Router ID: 10.255.70.103
                Primary Routing Table bgp.l2vpn.0

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.70.103:1
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:15:24   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Label-base: 800008, range: 8
                Localpref: 100
                Router ID: 10.255.70.103
                Secondary Tables: green.l2vpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.255.70.103 Metric: 2
                        Indirect next hop: 2 no-forward
                        Indirect path forwarding next hops: 1
                Next hop:        via so-0/3/0.0 weight 0x1
                        10.255.70.103/32 Originating RIB: inet.3
                          Metric: 2                       Node path count: 1
```

 Copyright © 2018, Juniper Networks, Inc.

```
                          Forwarding nexthops: 1
                                Nexthop: via so-0/3/0.0

10.255.70.103:2:3:1/96 (1 entry, 0 announced)
        *BGP    Preference: 170/-1
                Route Distinguisher: 10.255.70.103:2
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:15:24   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                AS path: I
                Communities: target:11111:2 Layer2-info: encaps:VPLS,
                control flags:Site-Down,
                mtu: 0
                Label-base: 800016, range: 8
                Localpref: 0
                Router ID: 10.255.70.103
                Secondary Tables: red.l2vpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.255.70.103 Metric: 2
                        Indirect next hop: 2 no-forward
                        Indirect path forwarding next hops: 1
                Next hop:        via so-0/3/0.0 weight 0x1
                        10.255.70.103/32 Originating RIB: inet.3
                          Metric: 2                       Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: via so-0/3/0.0
```

Copyright © 2018, Juniper Networks, Inc.

show route summary

List of Syntax

Syntax on page 2374
Syntax (EX Series Switches) on page 2374

Syntax

show route summary
<logical-system (all | *logical-system-name*)>
<table *routing-table-name*>

Syntax (EX Series Switches)

show route summary

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display summary statistics about the entries in the routing table.

CPU utilization might increase while the device learns routes. We recommend that you use the **show route summary** command after the device learns and enters the routes into the routing table. Depending on the size of your network, this might take several minutes. If you receive a “timeout communicating with routing daemon” error when using the **show route summary** command, wait several minutes before attempting to use the command again. This is not a critical system error, but you might experience a delay in using the command-line interface (CLI).

Options

none—Display summary statistics about the entries in the routing table.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

table *routing-table-name*—(Optional) Display summary statistics for all routing tables whose name begins with this string (for example, **inet.0** and **inet6.0** are both displayed when you run the **show route summary table inet** command). If you only want to display statistics for a specific routing table, make sure to enter the exact name of that routing table.

Required Privilege Level

view

List of Sample Output

show route summary on page 2376
show route summary table on page 2376
show route summary table (with Route Limits Configured for the Routing Table) on page 2377

Output Fields

Table 211 on page 2375 lists the output fields for the **show route summary** command. Output fields are listed in the approximate order in which they appear.

 Copyright © 2018, Juniper Networks, Inc.

Table 211: show route summary Output Fields

Field Name	Field Description
Router ID	Address of the local routing device.
routing-table-name	Name of the routing table (for example, **inet.0**).
destinations	Number of destinations for which there are routes in the routing table.
routes	Number of routes in the routing table: • **active**—Number of routes that are active. • **holddown**—Number of routes that are in the hold-down state before being declared inactive. • **hidden**—Number of routes that are not used because of routing policy.

Restart complete	All protocols have restarted for this routing table.
	Restart state:
	• **Pending:*protocol-name***—List of protocols that have not yet completed graceful restart for this routing table.
	• **Complete**—All protocols have restarted for this routing table.
	For example, if the output shows-
	• `LDP.inet.0: 5 routes (4 active, 1 holddown, 0 hidden)`
	`    Restart Pending: OSPF LDP VPN`
	This indicates that **OSPF**, **LDP**, and **VPN** protocols did not restart for **LDP.inet.0** routing table.
	• `vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)`
	`Restart Complete`
	This indicates that all protocols have restarted for **vpls_1.l2vpn.0** routing table.

Limit/Threshold	Displays the configured route limits for the routing table set with the ***maximum-prefixes*** and the **maximum-paths** statements. If you do not configure route limits for the routing table, the show output does not display this information. • **destinations**—The first number represents the maximum number of route prefixes installed in the routing table. The second number represents the number of route prefixes that trigger a warning message. • **routes**—The first number represents the maximum number of routes. The second number represents the number of routes that trigger a warning message.
Direct	Routes on the directly connected network.
Local	Local routes.
protocol-name	Name of the protocol from which the route was learned. For example, **OSPF**, **RSVP**, and **Static**.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show route summary

```
user@host> show route summary
Autonomous system number: 69
Router ID: 10.255.71.52
Maximum-ECMP: 32
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
              Direct:      6 routes,      5 active
               Local:      4 routes,      4 active
                OSPF:      5 routes,      4 active
              Static:      7 routes,      7 active
                IGMP:      1 routes,      1 active
                 PIM:      2 routes,      2 active

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
                RSVP:      2 routes,      2 active

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete
              Direct:      1 routes,      1 active

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
                MPLS:      3 routes,      3 active
                VPLS:      4 routes,      2 active

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete
              Direct:      2 routes,      2 active
                 PIM:      2 routes,      2 active
                 MLD:      1 routes,      1 active

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
                 BGP:      2 routes,      2 active
               L2VPN:      2 routes,      2 active

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete
                 BGP:      2 routes,      2 active
               L2VPN:      1 routes,      1 active

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
                 BGP:      4 routes,      4 active
```

show route summary table

```
user@host> show route summary table inet
Router ID: 192.168.0.1

inet.0: 32 destinations, 34 routes (31 active, 0 holddown, 1 hidden)
              Direct:      6 routes,      5 active
               Local:      9 routes,      9 active
                OSPF:      3 routes,      1 active
              Static:     13 routes,     13 active
```

 Copyright © 2018, Juniper Networks, Inc.

```
                IGMP:      1 routes,      1 active
                 PIM:      2 routes,      2 active

inet.1: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
           Multicast:      1 routes,      1 active

inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
               Local:      1 routes,      1 active
                 PIM:      2 routes,      2 active

inet6.1: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
           Multicast:      1 routes,      1 active
```

show route summary table (with Route Limits Configured for the Routing Table)

```
user@host> show route summary table VPN-A.inet.0
Autonomous system number: 100
Router ID: 10.255.182.142

VPN-A.inet.0: 13 destinations, 14 routes (13 active, 0 holddown, 0 hidden)
Limit/Threshold: 2000/200 destinations 20/12 routes
              Direct:      2 routes,      2 active
               Local:      1 routes,      1 active
                OSPF:      4 routes,      3 active
                 BGP:      4 routes,      4 active
                IGMP:      1 routes,      1 active
                 PIM:      2 routes,      2 active
```

Copyright © 2018, Juniper Networks, Inc.

show route table

List of Syntax

Syntax on page 2378
Syntax (EX Series Switches, QFX Series Switches) on page 2378

Syntax

show route table *routing-table-name*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches, QFX Series Switches)

show route table *routing-table-name*
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.
Show route table evpn statement introduced in Junos OS Release 15.1X53-D30 for QFX Series switches.

Description

Display the route entries in a particular routing table.

Options

brief | detail | extensive | terse—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

routing-table-name—Display route entries for all routing tables whose names begin with this string (for example, inet.0 and inet6.0 are both displayed when you run the **show route table inet** command).

Required Privilege Level

view

Related Documentation

- show route summary on page 2374

List of Sample Output

show route table bgp.l2.vpn on page 2389
show route table bgp.l3vpn.0 on page 2389
show route table bgp.l3vpn.0 detail on page 2390
show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured) on page 2391
show route table bgp.evpn.0 on page 2391
show route table evpna.evpn.0 on page 2392
show route table inet.0 on page 2392
show route table inet.3 on page 2393
show route table inet.3 protocol ospf on page 2393
show route table inet6.0 on page 2393
show route table inet6.3 on page 2393

Copyright © 2018, Juniper Networks, Inc.

Output Fields Table 193 on page 2170 describes the output fields for the **show route table** command. Output fields are listed in the approximate order in which they appear.

Table 212: show route table Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).

Restart complete	All protocols have restarted for this routing table.
	
	Restart state:
	
	• **Pending:***protocol-name*—List of protocols that have not yet completed graceful restart for this routing table.
	• **Complete**—All protocols have restarted for this routing table.
	
	For example, if the output shows-
	
	• `LDP.inet.0             : 5 routes (4 active, 1 holddown, 0 hidden)`
	`    Restart Pending: OSPF LDP VPN`
	
	This indicates that **OSPF**, **LDP**, and **VPN** protocols did not restart for the **LDP.inet.0** routing table.
	• `vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)`
	`Restart Complete`
	
	This indicates that all protocols have restarted for the **vpls_1.l2vpn.0** routing table.

number destinations	Number of destinations for which there are routes in the routing table.

Copyright © 2018, Juniper Networks, Inc.

Table 212: show route table Output Fields (continued)

Field Name	Field Description
number routes	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active) • **holddown** (routes that are in the pending state before being declared inactive) • **hidden** (routes that are not used because of a routing policy)

route-destination (entry, announced)	Route destination (for example:10.0.0.1/24). The **entry** value is the number of routes for this destination, and the **announced** value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:
	• ***MPLS-label*** (for example, 80001).
	• ***interface-name*** (for example, ge-1/0/2).
	• ***neighbor-address*:*control-word-status*:*encapsulation type*:*vc-id*:*source*** (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96).
	  ◦ ***neighbor-address***—Address of the neighbor.
	  ◦ ***control-word-status***—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**.
	  ◦ ***encapsulation type***—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.
	  ◦ ***vc-id***—Virtual circuit identifier.
	  ◦ ***source***—Source of the advertisement: **Local** or **Remote**.
	• ***inclusive multicast Ethernet tag route***—Type of route destination represented by (for example, 3:100.100.100.10:100::0::10::100.100.100.10/384):
	  ◦ *route distinguisher*—(8 octets) Route distinguisher (RD) must be the RD of the EVPN instance (EVI) that is advertising the NLRI.
	  ◦ *Ethernet tag ID*—(4 octets) Identifier of the Ethernet tag. Can set to 0 or to a valid Ethernet tag value.
	  ◦ *IP address length*—(1 octet) Length of IP address in bits.
	  ◦ *originating router's IP address*—(4 or 16 octets) Must set to the provider edge (PE) device's IP address. This address should be common for all EVIs on the PE device, and may be the PE device's loopback address.

label stacking	(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels. • **S=0 route** indicates that a packet with an incoming label stack depth of 2 or more exits this routing device with one fewer label (the label-popping operation is performed). • If there is no **S=** information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).

 Copyright © 2018, Juniper Networks, Inc.

Table 212: show route table Output Fields (continued)

Field Name	Field Description
[*protocol*, *preference*]	Protocol from which the route was learned and the preference value for the route. • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.
Level	(IS-IS only). In IS-IS, a single AS can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.
Route Distinguisher	IP subnet augmented with a 64-bit prefix.
PMSI	Provider multicast service interface (MVPN routing table).
Next-hop type	Type of next hop. For a description of possible values for this field, see Table 197 on page 2226.
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word **Selected**. This field can also contain the following information: • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.
Label-switched-path *lsp-path-name*	Name of the LSP used to reach the next hop.

Copyright © 2018, Juniper Networks, Inc.

Table 212: show route table Output Fields (continued)

Field Name	Field Description
Label operation	MPLS label and operation occurring at this routing device. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label).
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.
State	State of the route (a route can be in more than one state). See Table 198 on page 2228.
Local AS	AS number of the local routing devices.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.
Metric*n*	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances.
Task	Name of the protocol that has added the route.
Announcement bits	The number of BGP peers or protocols to which Junos OS has announced this route, followed by the list of the recipients of the announcement. Junos OS can also announce the route to the kernel routing table (KRT) for installing the route into the Packet Forwarding Engine, to a resolve tree, a Layer 2 VC, or even a VPN. For example, ***n*-Resolve inet** indicates that the specified route is used for route resolution for next hops found in the routing table. • *n*—An index used by Juniper Networks customer support only.

 Copyright © 2018, Juniper Networks, Inc.

Table 212: show route table Output Fields (continued)

Field Name	Field Description
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **Recorded**—The AS path is recorded by the sample process (sampled). • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used in the AS-path merge process, as defined in RFC 4893. • **[ ]**—If more than one AS number is configured on the routing device, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.
validation-state	(BGP-learned routes) Validation status of the route: • **Invalid**—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • **Unknown**—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • **Unverified**—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers. • **Valid**—Indicates that the prefix and autonomous system pair are found in the database.
FECs bound to route	Indicates point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
Primary Upstream	When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, indicates the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.
RPF Nexthops	When multipoint LDP with MoFRR is configured, indicates the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.
Label	Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.

Copyright © 2018, Juniper Networks, Inc.

Table 212: show route table Output Fields (continued)

Field Name	Field Description
weight	Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Prefixes bound to route	Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See Table 199 on page 2230 for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: **none** or **Site Down**.
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).
Accepted Multipath	Current active path when BGP multipath is configured.
Accepted LongLivedStale	The LongLivedStale flag indicates that the route was marked LLGR-stale by this router, as part of the operation of LLGR receiver mode. Either this flag or the LongLivedStaleImport flag might be displayed for a route. Neither of these flags is displayed at the same time as the Stale (ordinary GR stale) flag.
Accepted LongLivedStaleImport	The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy. Either this flag or the LongLivedStale flag might be displayed for a route. Neither of these flags is displayed at the same time as the Stale (ordinary GR stale) flag. Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and import into the inet.0 routing table
ImportAccepted LongLivedStaleImport	Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and imported into the inet.0 routing table The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy.
Accepted MultipathContrib	Path currently contributing to BGP multipath.
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.

 Copyright © 2018, Juniper Networks, Inc.

Table 212: show route table Output Fields (continued)

Field Name	Field Description
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.

Table 197 on page 2226 describes all possible values for the Next-hop Types output field.

Table 213: Next-hop Types Output Field Values

Next-Hop Type	Description
Broadcast (bcast)	Broadcast next hop.
Deny	Deny next hop.
Discard	Discard next hop.
Flood	Flood next hop. Consists of components called branches, up to a maximum of 32 branches. Each flood next-hop branch sends a copy of the traffic to the forwarding interface. Used by point-to-multipoint RSVP, point-to-multipoint LDP, point-to-multipoint CCC, and multicast.
Hold	Next hop is waiting to be resolved into a unicast or multicast type.
Indexed (idxd)	Indexed next hop.
Indirect (indr)	Used with applications that have a protocol next hop address that is remote. You are likely to see this next-hop type for internal BGP (IBGP) routes when the BGP next hop is a BGP neighbor that is not directly connected.
Interface	Used for a network address assigned to an interface. Unlike the router next hop, the interface next hop does not reference any specific node on the network.
Local (locl)	Local address on an interface. This next-hop type causes packets with this destination address to be received locally.
Multicast (mcst)	Wire multicast next hop (limited to the LAN).
Multicast discard (mdsc)	Multicast discard.
Multicast group (mgrp)	Multicast group member.
Receive (recv)	Receive.
Reject (rjct)	Discard. An ICMP unreachable message was sent.

Copyright © 2018, Juniper Networks, Inc.

Table 213: Next-hop Types Output Field Values (continued)

Next-Hop Type	Description
Resolve (rslv)	Resolving next hop.
Routed multicast (mcrt)	Regular multicast next hop.
Router	A specific node or set of nodes to which the routing device forwards packets that match the route prefix. To qualify as a next-hop type router, the route must meet the following criteria: • Must not be a direct or local subnet for the routing device. • Must have a next hop that is directly connected to the routing device.
Table	Routing table next hop.
Unicast (ucst)	Unicast.
Unilist (ulst)	List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.

Table 198 on page 2228 describes all possible values for the State output field. A route can be in more than one state (for example, **<Active NoReadvrt Int Ext>**).

Table 214: State Output Field Values

Value	Description
Accounting	Route needs accounting.
Active	Route is active.
Always Compare MED	Path with a lower multiple exit discriminator (MED) is available.
AS path	Shorter AS path is available.
Cisco Non-deterministic MED selection	Cisco nondeterministic MED is enabled, and a path with a lower MED is available.
Clone	Route is a clone.
Cluster list length	Length of cluster list sent by the route reflector.
Delete	Route has been deleted.
Ex	Exterior route.

Copyright © 2018, Juniper Networks, Inc.

Table 214: State Output Field Values (continued)

Value	Description
Ext	BGP route received from an external BGP neighbor.
FlashAll	Forces all protocols to be notified of a change to any route, active or inactive, for a prefix. When not set, protocols are informed of a prefix only when the active route changes.
Hidden	Route not used because of routing policy.
IfCheck	Route needs forwarding RPF check.
IGP metric	Path through next hop with lower IGP metric is available.
Inactive reason	Flags for this route, which was not selected as best for a particular destination.
Initial	Route being added.
Int	Interior route.
Int Ext	BGP route received from an internal BGP peer or a BGP confederation peer.
Interior > Exterior > Exterior via Interior	Direct, static, IGP, or EBGP path is available.
Local Preference	Path with a higher local preference value is available.
Martian	Route is a martian (ignored because it is obviously invalid).
MartianOK	Route exempt from martian filtering.
Next hop address	Path with lower metric next hop is available.
No difference	Path from neighbor with lower IP address is available.
NoReadvrt	Route not to be advertised.
NotBest	Route not chosen because it does not have the lowest MED.
Not Best in its group	Incoming BGP AS is not the best of a group (only one AS can be the best).
NotInstall	Route not to be installed in the forwarding table.
Number of gateways	Path with a greater number of next hops is available.
Origin	Path with a lower origin code is available.

Copyright © 2018, Juniper Networks, Inc.

Table 214: State Output Field Values (continued)

Value	Description
Pending	Route pending because of a hold-down configured on another route.
Release	Route scheduled for release.
RIB preference	Route from a higher-numbered routing table is available.
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.
Route Metric or MED comparison	Route with a lower metric or MED is available.
Route Preference	Route with lower preference value is available.
Router ID	Path through a neighbor with lower ID is available.
Secondary	Route not a primary route.
Unusable path	Path is not usable because of one of the following conditions: • The route is damped. • The route is rejected by an import policy. • The route is unresolved.
Update source	Last tiebreaker is the lowest IP address value.

Table 199 on page 2230 describes the possible values for the Communities output field.

Table 215: Communities Output Field Values

Value	Description
area-number	4 bytes, encoding a 32-bit area number. For AS-external routes, the value is **0**. A nonzero value identifies the route as internal to the OSPF domain, and as within the identified area. Area numbers are relative to a particular OSPF domain.
bandwidth: ***local AS number:link-bandwidth-number***	Link-bandwidth community value used for unequal-cost load balancing. When BGP has several candidate paths available for multipath purposes, it does not perform unequal-cost load balancing according to the link-bandwidth community unless all candidate paths have this attribute.
domain-id	Unique configurable number that identifies the OSPF domain.
domain-id-vendor	Unique configurable number that further identifies the OSPF domain.
link-bandwidth-number	Link-bandwidth number: from **0** through **4,294,967,295** (bytes per second).
local AS number	Local AS number: from **1** through **65,535**.

Copyright © 2018, Juniper Networks, Inc.

Table 215: Communities Output Field Values (continued)

Value	Description
options	1 byte. Currently this is only used if the route type is **5** or **7**. Setting the least significant bit in the field indicates that the route carries a type 2 metric.
origin	(Used with VPNs) Identifies where the route came from.
ospf-route-type	1 byte, encoded as **1** or **2** for intra-area routes (depending on whether the route came from a type 1 or a type 2 LSA); **3** for summary routes; **5** for external routes (area number must be**0**); **7** for NSSA routes; or **129** for sham link endpoint addresses.
route-type-vendor	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute **0x8000**. The format is ***area-number:ospf-route-type:options***.
rte-type	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute **0x0306**. The format is ***area-number:ospf-route-type:options***.
target	Defines which VPN the route participates in; **target** has the format ***32-bit IP address:16-bit number***. For example, 10.19.0.0:100.
unknown IANA	Incoming IANA codes with a value between **0x1** and **0x7fff**. This code of the BGP extended community attribute is accepted, but it is not recognized.
unknown OSPF vendor community	Incoming IANA codes with a value above **0x8000**. This code of the BGP extended community attribute is accepted, but it is not recognized.

Sample Output

show route table bgp.l2.vpn

```
user@host> show route table bgp.l2.vpn
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.24.1:1:4:1/96
                   *[BGP/170] 01:08:58, localpref 100, from 192.168.24.1
                      AS path: I
                    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

show route table bgp.l3vpn.0

```
user@host> show route table bgp.l3vpn.0
      bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
      + = Active Route, - = Last Active, * = Both

      10.255.71.15:100:10.255.71.17/32
                         *[BGP/170] 00:03:59, MED 1, localpref 100, from
      10.255.71.15
                            AS path: I
                          > via so-2/1/0.0, Push 100020, Push 100011(top)
      10.255.71.15:200:10.255.71.18/32
```

```
                        *[BGP/170] 00:03:59, MED 1, localpref 100, from
10.255.71.15
                          AS path: I
                        > via so-2/1/0.0, Push 100021, Push 100011(top)
```

show route table bgp.l3vpn.0 detail

```
user@host> show route table bgp.l3vpn.0 detail
bgp.l3vpn.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)

10.255.245.12:1:172.16.4.0/8 (1 entry, 1 announced)
  *BGP    Preference: 170/-101
          Route Distinguisher: 10.255.245.12:1
          Source: 10.255.245.12
          Next hop: 192.168.208.66 via fe-0/0/0.0, selected
          Label operation: Push 182449
          Protocol next hop: 10.255.245.12
          Push 182449
          Indirect next hop: 863a630 297
          State: <Active Int Ext>
          Local AS:    35 Peer AS:    35
          Age: 12:19      Metric2: 1
          Task: BGP_35.10.255.245.12+179
          Announcement bits (1): 0-BGP.0.0.0.0+179
          AS path: 30 10458 14203 2914 3356 I (Atomic) Aggregator: 3356 4.68.0.11

          Communities: 2914:420 target:11111:1 origin:56:78
          VPN Label: 182449
          Localpref: 100
           Router ID: 10.255.245.12

10.255.245.12:1:4.17.225.0/24 (1 entry, 1 announced)
  *BGP    Preference: 170/-101
          Route Distinguisher: 10.255.245.12:1
          Source: 10.255.245.12
          Next hop: 192.168.208.66 via fe-0/0/0.0, selected
          Label operation: Push 182465
          Protocol next hop: 10.255.245.12
          Push 182465
          Indirect next hop: 863a8f0 305
          State: <Active Int Ext>
          Local AS:    35 Peer AS:    35
          Age: 12:19      Metric2: 1
          Task: BGP_35.10.255.245.12+179
          Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496 6496 I
                Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
                VPN Label: 182465
                Localpref: 100
                Router ID: 10.255.245.12

10.255.245.12:1:4.17.226.0/23 (1 entry, 1 announced)
  *BGP    Preference: 170/-101
          Route Distinguisher: 10.255.245.12:1
          Source: 10.255.245.12
          Next hop: 192.168.208.66 via fe-0/0/0.0, selected
          Label operation: Push 182465
          Protocol next hop: 10.255.245.12
          Push 182465
          Indirect next hop: 86bd210 330
          State: <Active Int Ext>
```


Copyright © 2018, Juniper Networks, Inc.

```
                Local AS:    35 Peer AS:    35
                Age: 12:19      Metric2: 1
                Task: BGP_35.10.255.245.12+179
                Announcement bits (1): 0-BGP.0.0.0.0+179
                AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496

                6496 I
                Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
                VPN Label: 182465
                Localpref: 100
                Router ID: 10.255.245.12

10.255.245.12:1:4.17.251.0/24 (1 entry, 1 announced)
  *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.245.12:1
                Source: 10.255.245.12
                Next hop: 192.168.208.66 via fe-0/0/0.0, selected
                Label operation: Push 182465
                Protocol next hop: 10.255.245.12
                Push 182465
                Indirect next hop: 86bd210 330
                State: <Active Int Ext>
                Local AS:    35 Peer AS:    35
                Age: 12:19      Metric2: 1
                Task: BGP_35.10.255.245.12+179
                Announcement bits (1): 0-BGP.0.0.0.0+179
                AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496

                6496 I
                Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
                VPN Label: 182465
                Localpref: 100
```

show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured)

```
user@host> show route table bgp.rtarget.o
bgp.rtarget.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100:100:100/96
                   *[RTarget/5] 00:03:14
                      Type Proxy
                        for 10.255.165.103
                        for 10.255.166.124
                      Local
```

show route table bgp.evpn.0

```
user@host> show route table bgp.evpn.0
bgp.evpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2:100.100.100.2:100::0::00:26:88:5f:67:b0/304
                   *[BGP/170] 11:00:05, localpref 100, from 100.100.100.2
                      AS path: I, validation-state: unverified
                    > to 100.64.12.2 via xe-2/2/0.0, label-switched-path R0toR1
2:100.100.100.2:100::0::00:51:51:51:51:51/304
                   *[BGP/170] 11:00:05, localpref 100, from 100.100.100.2
                      AS path: I, validation-state: unverified
                    > to 100.64.12.2 via xe-2/2/0.0, label-switched-path R0toR1
```

```
2:100.100.100.3:100::0::00:52:52:52:52:52/304
                   *[BGP/170] 10:59:58, localpref 100, from 100.100.100.3
                      AS path: I, validation-state: unverified
                    > to 100.64.13.3 via ge-2/0/8.0, label-switched-path R0toR2
2:100.100.100.3:100::0::a8:d0:e5:5b:01:c8/304
                   *[BGP/170] 10:59:58, localpref 100, from 100.100.100.3
                      AS path: I, validation-state: unverified
                    > to 100.64.13.3 via ge-2/0/8.0, label-switched-path R0toR2
3:100.100.100.2:100::1000::100.100.100.2/304
                   *[BGP/170] 11:00:16, localpref 100, from 100.100.100.2
                      AS path: I, validation-state: unverified
                    > to 100.64.12.2 via xe-2/2/0.0, label-switched-path R0toR1
3:100.100.100.2:100::2000::100.100.100.2/304
                   *[BGP/170] 11:00:16, localpref 100, from 100.100.100.2
                      AS path: I, validation-state: unverified
                    > to 100.64.12.2 via xe-2/2/0.0, label-switched-path R0toR1
```

show route table evpna.evpn.0

```
user@host> show route table evpna.evpn.0
evpna.evpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

3:100.100.100.10:100::0::10::100.100.100.10/384
                   *[EVPN/170] 01:37:09
                      Indirect
3:100.100.100.2:100::2000::100.100.100.2/304
                   *[EVPN/170] 01:37:12
                      Indirect
```

show route table inet.0

```
user@host> show route table inet.0
inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 00:51:57
                    > to 172.16.5.254 via fxp0.0
10.0.0.1/32         *[Direct/0] 00:51:58
                    > via at-5/3/0.0
10.0.0.2/32         *[Local/0] 00:51:58
                     Local
10.12.12.21/32     *[Local/0] 00:51:57
                     Reject
10.13.13.13/32     *[Direct/0] 00:51:58
                    > via t3-5/2/1.0
10.13.13.14/32     *[Local/0] 00:51:58
                     Local
10.13.13.21/32     *[Local/0] 00:51:58
                     Local
10.13.13.22/32     *[Direct/0] 00:33:59
                    > via t3-5/2/0.0
127.0.0.1/32        [Direct/0] 00:51:58
                    > via lo0.0
10.222.5.0/24     *[Direct/0] 00:51:58
                    > via fxp0.0
10.222.5.81/32    *[Local/0] 00:51:58
                     Local
```


Copyright © 2018, Juniper Networks, Inc.

show route table inet.3

```
user@host> show route table inet.3
inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.5/32        *[LDP/9] 00:25:43, metric 10, tag 200
                      to 10.2.94.2 via lt-1/2/0.49
                    > to 10.2.3.2 via lt-1/2/0.23
```

show route table inet.3 protocol ospf

```
user@host> show route table inet.3 protocol ospf
inet.3: 9 destinations, 18 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.20/32         [L-OSPF/10] 1d 00:00:56, metric 2
                    > to 10.0.10.70 via lt-1/2/0.14, Push 800020
                     to 10.0.6.60 via lt-1/2/0.12, Push 800020, Push 800030(top)
1.1.1.30/32         [L-OSPF/10] 1d 00:01:01, metric 3
                    > to 10.0.10.70 via lt-1/2/0.14, Push 800030
                      to 10.0.6.60 via lt-1/2/0.12, Push 800030
1.1.1.40/32         [L-OSPF/10] 1d 00:01:01, metric 4
                    > to 10.0.10.70 via lt-1/2/0.14, Push 800040
                      to 10.0.6.60 via lt-1/2/0.12, Push 800040
1.1.1.50/32         [L-OSPF/10] 1d 00:01:01, metric 5
                    > to 10.0.10.70 via lt-1/2/0.14, Push 800050
                      to 10.0.6.60 via lt-1/2/0.12, Push 800050
1.1.1.60/32         [L-OSPF/10] 1d 00:01:01, metric 6
                    > to 10.0.10.70 via lt-1/2/0.14, Push 800060
                      to 10.0.6.60 via lt-1/2/0.12, Pop
```

show route table inet6.0

```
user@host> show route table inet6.0
inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Route, * = Both

fec0:0:0:3::/64 *[Direct/0] 00:01:34
>via fe-0/1/0.0

fec0:0:0:3::/128 *[Local/0] 00:01:34
>Local

fec0:0:0:4::/64 *[Static/5] 00:01:34
>to fec0:0:0:3::ffff via fe-0/1/0.0
```

show route table inet6.3

```
user@router> show route table inet6.3
inet6.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::10.255.245.195/128
                   *[LDP/9] 00:00:22, metric 1
                    > via so-1/0/0.0
::10.255.245.196/128
                   *[LDP/9] 00:00:08, metric 1
                    > via so-1/0/0.0, Push 100008
```

Copyright © 2018, Juniper Networks, Inc.

show route table inetflow detail

```
user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Next-hop reference count: 2
                State: <Active Ext>
                Local AS: 64502 Peer AS: 64500
                Age: 4
                Task: BGP_64500.10.12.99.5+3792
                Announcement bits (1): 0-Flow
                AS path: 64500 I
                Communities: traffic-rate:0:0
                Validation state: Accept, Originator: 10.12.99.5
                Via: 10.12.44.0/24, Active
                Localpref: 100
                Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
        *Flow   Preference: 5
                Next-hop reference count: 2
                State: <Active>
                Local AS: 64502
                Age: 6:30
                Task: RT Flow
                Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
                AS path: I
                Communities: 1:1
```

show route table lsdist.0 extensive

```
user@host> show route table lsdist.0 extensive
lsdist.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
NODE { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 ISIS-L1:0 }/1152
 (1 entry, 1 announced)
TSI:
Page 0 idx 0, (group ibgp type Internal) Type 1 val 0xa62f378 (adv_entry)
   Advertised metrics:
     Nexthop: Self
     Localpref: 100
     AS path: [4170512532] I
     Communities:
Path NODE { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 ISIS-L1:0 }
 Vector len 4.  Val: 0
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Fictitious, Next hop index: 0
                Address: 0x95dfc64
                Next-hop reference count: 9
                State: <Active NotInstall>
                Local AS: 4170512532
                Age: 6:05
                Validation State: unverified
                Task: IS-IS
                Announcement bits (1): 0-BGP_RT_Background
                AS path: I
                IPv4 Router-ids:
                   128.220.11.197
                Area membership:
```

 Copyright © 2018, Juniper Networks, Inc.

```
                  47 00 05 80 ff f8 00 00 00 01 08 00 01
              SPRING-Capabilities:                 - SRGB block [Start: 800000,
 Range: 256, Flags: 0xc0]
                SPRING-Algorithms:
                   - Algo: 0
 LINK { Local { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 }.{
IPv4:8.65.1.105 } Remote { AS:4170512532 BGP-LS ID:4170512532 ISO:4284.3300.5067)
TSI:
Page 0 idx 0, (group ibgp type Internal) Type 1 val 0xa62f3cc (adv_entry)
    Advertised metrics:
      Nexthop: Self
      Localpref: 100
      AS path: [4170512532] I
      Communities:
Path LINK { Local { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 }.{
 IPv4:8.65.1.105 } Remote { AS:4170512532 BGP-LS ID:4170512532 ISO:4284.33000
      *IS-IS  Preference: 15
                Level: 1
                Next hop type: Fictitious, Next hop index: 0
                Address: 0x95dfc64
                Next-hop reference count: 9
                State: <Active NotInstall>
                Local AS: 4170512532
                Age: 6:05
                Validation State: unverified
                Task: IS-IS
                Announcement bits (1): 0-BGP_RT_Background
                AS path: I
                Color: 32768
                Maximum bandwidth: 1000Mbps
                Reservable bandwidth: 1000Mbps
                Unreserved bandwidth by priority:
                   0   1000Mbps
                   1   1000Mbps
                   2   1000Mbps
                   3   1000Mbps
                   4   1000Mbps
                   5   1000Mbps
                   6   1000Mbps
                   7   1000Mbps
                 Metric: 10
                 TE Metric: 10
                 LAN IPV4 Adj-SID -  Label: 299776, Flags: 0x30,
Weight: 0, Nbr: 10.220.1.83

PREFIX { Node { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 } {
IPv4:128.220.11.197/32 } ISIS-L1:0 }/1152 (1 entry, 1 announced) TSI: Page 0 idx
 0, (group ibgp type Internal) Type 1 val 0xa62f43c (adv_entry)
    Advertised metrics:
      Nexthop: Self
      Localpref: 100
      AS path: [4170512532] I
      Communities:
Path PREFIX { Node { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 }
{ IPv4:128.220.11.197/32 } ISIS-L1:0 } Vector len 4.  Val: 0
         *IS-IS  Preference: 15
                 Level: 1
                 Next hop type: Fictitious, Next hop index: 0
                 Address: 0x95dfc64
                 Next-hop reference count: 9
                 State:<Active NotInstall>
```

```
                Local AS: 4170512532
                Age: 6:05
                Validation State: unverified
                Task: IS-IS
                Announcement bits (1): 0-BGP_RT_Background
                AS path: I
                              Prefix SID: 67, Flags: 0x40, Algo: 0
```

show route table l2circuit.0

```
user@host> show route table l2circuit.0
l2circuit.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.195:NoCtrlWord:1:1:Local/96
                   *[L2CKT/7] 00:50:47
                    > via so-0/1/2.0, Push 100049
                      via so-0/1/3.0, Push 100049
10.1.1.195:NoCtrlWord:1:1:Remote/96
                   *[LDP/9] 00:50:14
                      Discard
10.1.1.195:CtrlWord:1:2:Local/96
                   *[L2CKT/7] 00:50:47
                    > via so-0/1/2.0, Push 100049
                      via so-0/1/3.0, Push 100049
10.1.1.195:CtrlWord:1:2:Remote/96
                   *[LDP/9] 00:50:14
                      Discard
```

show route table mpls

```
user@host> show route table mpls
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 00:13:55, metric 1
                      Receive
1                  *[MPLS/0] 00:13:55, metric 1
                      Receive
2                  *[MPLS/0] 00:13:55, metric 1
                      Receive
1024               *[VPN/0] 00:04:18
                      to table red.inet.0, Pop
```

show route table mpls extensive

```
user@host> show route table mpls extensive
100000 (1 entry, 1 announced)
TSI:
KRT in-kernel 100000 /36 -> {so-1/0/0.0}
        *LDP    Preference:   9
                Next hop:  via so-1/0/0.0, selected
                Pop
                State: <Active Int>
                Age: 29:50      Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                Prefixes bound to route: 10.0.0.194/32
```

 Copyright © 2018, Juniper Networks, Inc.

show route table mpls.0

```
user@host> show route table mpls.0
mpls.0: 18 destinations, 19 routes (18 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 11:39:56, metric 1
                      to table inet.0
0(S=0)             *[MPLS/0] 11:39:56, metric 1
                      to table mpls.0
1                  *[MPLS/0] 11:39:56, metric 1
                      Receive
2                  *[MPLS/0] 11:39:56, metric 1
                      to table inet6.0
2(S=0)             *[MPLS/0] 11:39:56, metric 1
                      to table mpls.0
13                 *[MPLS/0] 11:39:56, metric 1
                      Receive
303168             *[EVPN/7] 11:00:49, routing-instance pbbn10, route-type
Ingress-MAC, ISID 0
                      to table pbbn10.evpn-mac.0
303184             *[EVPN/7] 11:00:53, routing-instance pbbn10, route-type
Ingress-IM, ISID 1000
                      to table pbbn10.evpn-mac.0
                    [EVPN/7] 11:00:53, routing-instance pbbn10, route-type
Ingress-IM, ISID 2000
                      to table pbbn10.evpn-mac.0
303264             *[EVPN/7] 11:00:53, remote-pe 100.100.100.2, routing-instance
 pbbn10, route-type Egress-IM, ISID 1000
                    > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303280             *[EVPN/7] 11:00:53, remote-pe 100.100.100.2, routing-instance
 pbbn10, route-type Egress-IM, ISID 2000
                    > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303328             *[EVPN/7] 11:00:49, remote-pe 100.100.100.2, routing-instance
 pbbn10, route-type Egress-MAC, ISID 0
                    > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303344             *[EVPN/7] 11:00:49, remote-pe 100.100.100.2, routing-instance
 pbbn10, route-type Egress-MAC, ISID 0
                    > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303360             *[EVPN/7] 11:00:47, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC 00:26:88:5f:67:b0
                    > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303376             *[EVPN/7] 11:00:47, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC 00:51:51:51:51:51
                    > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303392             *[EVPN/7] 11:00:35, remote-pe 100.100.100.3, routing-instance
 pbbn10, route-type Egress-MAC, ISID 0
                    > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
303408             *[EVPN/7] 11:00:35, remote-pe 100.100.100.3, routing-instance
 pbbn10, route-type Egress-MAC, ISID 0
                    > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
303424             *[EVPN/7] 11:00:33, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC a8:d0:e5:5b:01:c8
                    > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
303440             *[EVPN/7] 11:00:33, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC 00:52:52:52:52:52
                    > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
```

show route table mpls.0 detail (PTX Series)

```
user@host> show route table mpls.0 detail
ge-0/0/2.600 (1 entry, 1 announced)
        *L2VPN  Preference: 7
                Next hop type: Indirect
                Address: 0x9438f34
                Next-hop reference count: 2
                Next hop type: Router, Next hop index: 567
                Next hop: 10.0.0.1 via ge-0/0/1.0, selected
                Label operation: Push 299808
                Label TTL action: prop-ttl
                Load balance label: Label 299808:None;
                Session Id: 0x1
                Protocol next hop: 10.255.255.1
                Label operation: Push 299872 Offset: 252
                Label TTL action: no-prop-ttl
                Load balance label: Label 299872:Flow label PUSH;
                Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
                Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
                State: <Active Int>
                Age: 21         Metric2: 1
                Validation State: unverified
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 2-Common L2 VC
                AS path: I
```

show route table mpls.0 ccc ge-0/0/1.1004 detail

```
user@host>show route table mpls.0 ccc ge-0/0/1.1004 detail
mpls.0: 121 destinations, 121 routes (121 active, 0 holddown, 0 hidden)
ge-0/0/1.1004 (1 entry, 1 announced)
        *EVPN   Preference: 7
                Next hop type: List, Next hop index: 1048577
                Address: 0xdc14770
                Next-hop reference count: 3
                Next hop: ELNH Address 0xd011e30
                    Next hop type: Indirect, Next hop index: 0
                    Address: 0xd011e30
                    Next-hop reference count: 3
                    Protocol next hop: 100.100.100.1
                    Label operation: Push 301952
                    Composite next hop: 0xd011dc0 754 INH Session ID: 0x146
                    Indirect next hop: 0xb69a890 1048615 INH Session ID: 0x146
                        Next hop type: Router, Next hop index: 735
                        Address: 0xd00e530
                        Next-hop reference count: 23
                        Next hop: 100.46.1.2 via ge-0/0/5.0
                        Label-switched-path pe4_to_pe1
                        Label operation: Push 300320
                        Label TTL action: prop-ttl
                        Load balance label: Label 300320: None;
                        Label element ptr: 0xd00e580
                        Label parent element ptr: 0x0
                        Label element references: 18
                        Label element child references: 16
                        Label element lsp id: 5
                Next hop: ELNH Address 0xd012070
                    Next hop type: Indirect, Next hop index: 0
                    Address: 0xd012070
```


Copyright © 2018, Juniper Networks, Inc.

```
                Next-hop reference count: 3
                Protocol next hop: 100.100.100.2
                Label operation: Push 301888
                Composite next hop: 0xd012000 755 INH Session ID: 0x143
                Indirect next hop: 0xb69a9a0 1048641 INH Session ID: 0x143
                    Next hop type: Router, Next hop index: 716
                    Address: 0xd00e710
                    Next-hop reference count: 23
                    Next hop: 100.46.1.2 via ge-0/0/5.0
                    Label-switched-path pe4_to_pe2
                    Label operation: Push 300304
                    Label TTL action: prop-ttl
                    Load balance label: Label 300304: None;
                    Label element ptr: 0xd00e760
                    Label parent element ptr: 0x0
                    Label element references: 15
                    Label element child references: 13
                    Label element lsp id: 6
            Next hop: ELNH Address 0xd0121f0, selected
                Next hop type: Indirect, Next hop index: 0
                Address: 0xd0121f0
                Next-hop reference count: 3
                Protocol next hop: 100.100.100.3
                Label operation: Push 301984
                Composite next hop: 0xd012180 756 INH Session ID: 0x145
                Indirect next hop: 0xb69aab0 1048642 INH Session ID: 0x145
                    Next hop type: Router, Next hop index: 801
                    Address: 0xd010ed0
                    Next-hop reference count: 32
                    Next hop: 100.46.1.2 via ge-0/0/5.0
                    Label-switched-path pe4_to_pe3
                    Label operation: Push 300336
                    Label TTL action: prop-ttl
                    Load balance label: Label 300336: None;
                    Label element ptr: 0xd0108c0
                    Label parent element ptr: 0x0
                    Label element references: 22
                    Label element child references: 20
                    Label element lsp id: 7
            State: < Active Int >
            Age: 2:06:50
            Validation State: unverified
            Task: evpn global task
            Announcement bits (1): 1-KRT
            AS path: I
```

show route table mpls.0 protocol evpn

```
user@host>show route table mpls.0 protocol evpn
mpls.0: 121 destinations, 121 routes (121 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

299872             *[EVPN/7] 02:30:58, routing-instance mhevpn, route-type
Ingress-IM, vlan-id 10
                      to table mhevpn.evpn-mac.0
300016             *[EVPN/7] 02:30:38, routing-instance VS-1, route-type
Ingress-IM, vlan-id 110
                      to table VS-1.evpn-mac.0
300032             *[EVPN/7] 02:30:38, routing-instance VS-1, route-type
Ingress-IM, vlan-id 120
                      to table VS-1.evpn-mac.0
```

```
300048              *[EVPN/7] 02:30:38, routing-instance VS-1, route-type
Ingress-IM, vlan-id 130
                       to table VS-1.evpn-mac.0
300064              *[EVPN/7] 02:30:38, routing-instance VS-2, route-type
Ingress-IM, vlan-id 210
                       to table VS-2.evpn-mac.0
300080              *[EVPN/7] 02:30:38, routing-instance VS-2, route-type
Ingress-IM, vlan-id 220
                       to table VS-2.evpn-mac.0
300096              *[EVPN/7] 02:30:38, routing-instance VS-2, route-type
Ingress-IM, vlan-id 230
                       to table VS-2.evpn-mac.0
300112              *[EVPN/7] 02:27:06, routing-instance mhevpn, route-type
Egress-MAC, ESI 00:44:44:44:44:44:44:44:44:44
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300128              *[EVPN/7] 02:29:22, routing-instance mhevpn, route-type
Ingress-Aliasing
                       to table mhevpn.evpn-mac.0
300144              *[EVPN/7] 02:27:06, routing-instance VS-1, route-type
Egress-MAC, ESI 00:44:44:44:44:44:44:44:44:44
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300160              *[EVPN/7] 02:29:22, routing-instance VS-1, route-type
Ingress-Aliasing
                       to table VS-1.evpn-mac.0
300176              *[EVPN/7] 02:27:07, routing-instance VS-2, route-type
Egress-MAC, ESI 00:44:44:44:44:44:44:44:44:44
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300192              *[EVPN/7] 02:29:22, routing-instance VS-2, route-type
Ingress-Aliasing
                       to table VS-2.evpn-mac.0
300208              *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-1, route-type Egress-IM, vlan-id 120
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300224              *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 mhevpn, route-type Egress-IM, vlan-id 10
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300240              *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-1, route-type Egress-IM, vlan-id 110
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300256              *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-1, route-type Egress-IM, vlan-id 130
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300272              *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-2, route-type Egress-IM, vlan-id 210
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300288              *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-2, route-type Egress-IM, vlan-id 220
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300304              *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-2, route-type Egress-IM, vlan-id 230
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300320              *[EVPN/7] 02:27:06, routing-instance VS-1, route-type
Egress-MAC, ESI 00:11:11:11:11:11:11:11:11:11
                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300336              *[EVPN/7] 02:27:06, routing-instance VS-1, route-type
Egress-MAC, ESI 00:33:33:33:33:33:33:33:33:33
                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
```

 Copyright © 2018, Juniper Networks, Inc.

```
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300368               *[EVPN/7] 02:27:07, routing-instance VS-2, route-type
Egress-MAC, ESI 00:33:33:33:33:33:33:33:33:33
                      to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300384               *[EVPN/7] 02:27:07, routing-instance VS-2, route-type
Egress-MAC, ESI 00:11:11:11:11:11:11:11:11:11
                      to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                      to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300416               *[EVPN/7] 02:27:06, routing-instance mhevpn, route-type
Egress-MAC, ESI 00:33:33:33:33:33:33:33:33:33
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                      to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300432               *[EVPN/7] 02:27:06, routing-instance mhevpn, route-type
Egress-MAC, ESI 00:11:11:11:11:11:11:11:11:11
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                      to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

                      to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300480               *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-1, route-type Egress-MAC
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300496               *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-2, route-type Egress-MAC
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300560               *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-1, route-type Egress-MAC
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300592               *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-2, route-type Egress-MAC
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300608               *[EVPN/7] 02:29:23
                     > via ge-0/0/1.1001, Pop
300624               *[EVPN/7] 02:29:23
                     > via ge-0/0/1.2001, Pop
301232               *[EVPN/7] 02:29:17
                     > via ge-0/0/1.1002, Pop
301296               *[EVPN/7] 02:29:10
                     > via ge-0/0/1.1003, Pop
301312               *[EVPN/7] 02:27:06
                     > via ae10.2003, Pop
                      to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301360               *[EVPN/7] 02:29:01
                     > via ge-0/0/1.1004, Pop
301408               *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 vpws1004, route-type Egress, vlan-id 2004
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
301456               *[EVPN/7] 02:27:06
                     > via ae10.1010, Pop
                      to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301552               *[EVPN/7] 02:27:07, routing-instance VS-1, route-type
Egress-MAC, ESI 00:22:22:22:22:22:22:22:22:22
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301568               *[EVPN/7] 02:27:07, routing-instance VS-2, route-type
```

```
Egress-MAC, ESI 00:22:22:22:22:22:22:22:22:22
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301648              *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 vpws1010, route-type Egress, vlan-id 2010
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
301664              *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 mhevpn, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
301680              *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 mhevpn, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
301696              *[EVPN/7] 02:27:07, routing-instance mhevpn, route-type
Egress-MAC, ESI 00:22:22:22:22:22:22:22:22:22
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301712              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-2, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301728              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-1, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301744              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-2, route-type Egress-IM, vlan-id 230
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301760              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 vpws1010, route-type Egress, vlan-id 2010
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301776              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 mhevpn, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301792              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-1, route-type Egress-IM, vlan-id 130
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301808              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 vpws1004, route-type Egress, vlan-id 2004
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301824              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 mhevpn, route-type Egress-IM, vlan-id 10
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301840              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 vpws1002, route-type Egress, vlan-id 2002
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301856              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 vpws1003, route-type Egress, vlan-id 2003
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301872              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 vpws1003, route-type Egress Protection, vlan-id 2003
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301888              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 vpws1010, route-type Egress Protection, vlan-id 1010
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301904              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-2, route-type Egress-IM, vlan-id 220
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301920              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-2, route-type Egress-IM, vlan-id 210
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301936              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-2, route-type Egress-IM, vlan-id 230
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301952              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-2, route-type Egress-SH, vlan-id 230
```

 Copyright © 2018, Juniper Networks, Inc.

```
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301968              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-2, route-type Egress-IM, vlan-id 220
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301984              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-2, route-type Egress-SH, vlan-id 220
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302000              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-2, route-type Egress-IM, vlan-id 210
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302016              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-2, route-type Egress-SH, vlan-id 210
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302032              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-2, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302048              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-2, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302064              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-2, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302080              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-2, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302096              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-1, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302112              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-1, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302128              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-1, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302144              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-1, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302160              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-1, route-type Egress-IM, vlan-id 120
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302176              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-1, route-type Egress-IM, vlan-id 110
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302192              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-1, route-type Egress-IM, vlan-id 130
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302208              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-1, route-type Egress-SH, vlan-id 130
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302224              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-1, route-type Egress-IM, vlan-id 120
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302240              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-1, route-type Egress-SH, vlan-id 120
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302256              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-1, route-type Egress-IM, vlan-id 110
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302272              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-1, route-type Egress-SH, vlan-id 110
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
```

Copyright © 2018, Juniper Networks, Inc.

```
302288             *[EVPN/7] 02:27:06, remote-pe 100.100.100.1, routing-instance
 mhevpn, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302304             *[EVPN/7] 02:27:06, remote-pe 100.100.100.1, routing-instance
 mhevpn, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302320             *[EVPN/7] 02:27:06, remote-pe 100.100.100.3, routing-instance
 mhevpn, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302336             *[EVPN/7] 02:27:06, remote-pe 100.100.100.3, routing-instance
 mhevpn, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302352             *[EVPN/7] 02:27:06, remote-pe 100.100.100.3, routing-instance
 vpws1004, route-type Egress, vlan-id 2004
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302368             *[EVPN/7] 02:27:06, remote-pe 100.100.100.3, routing-instance
 mhevpn, route-type Egress-IM, vlan-id 10
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302384             *[EVPN/7] 02:27:06, remote-pe 100.100.100.3, routing-instance
 mhevpn, route-type Egress-SH, vlan-id 10
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302400             *[EVPN/7] 02:26:21
                     > via ge-0/0/1.3001, Pop
302432             *[EVPN/7] 02:26:21, remote-pe 100.100.100.3, routing-instance
 vpws3001, route-type Egress, vlan-id 40000
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302448             *[EVPN/7] 02:26:21, remote-pe 100.100.100.1, routing-instance
 vpws3001, route-type Egress, vlan-id 40000
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302464             *[EVPN/7] 02:26:20, remote-pe 100.100.100.2, routing-instance
 vpws3001, route-type Egress, vlan-id 40000
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
302480             *[EVPN/7] 02:26:14
                     > via ge-0/0/1.3016, Pop
302512             *[EVPN/7] 02:26:14, remote-pe 100.100.100.1, routing-instance
 vpws3016, route-type Egress, vlan-id 40016
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302528             *[EVPN/7] 02:26:14, remote-pe 100.100.100.2, routing-instance
 vpws3016, route-type Egress, vlan-id 40016
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
302560             *[EVPN/7] 02:26:06
                     > via ae10.3011, Pop
                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302592             *[EVPN/7] 02:26:07, remote-pe 100.100.100.1, routing-instance
 vpws3011, route-type Egress, vlan-id 401100
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302608             *[EVPN/7] 02:26:07, remote-pe 100.100.100.2, routing-instance
 vpws3011, route-type Egress, vlan-id 401100
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
302624             *[EVPN/7] 02:26:07, remote-pe 100.100.100.3, routing-instance
 vpws3011, route-type Egress Protection, vlan-id 301100
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302656             *[EVPN/7] 02:25:59
                     > via ae10.3006, Pop
                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302688             *[EVPN/7] 02:26:00, remote-pe 100.100.100.2, routing-instance
 vpws3006, route-type Egress, vlan-id 400600
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
302704             *[EVPN/7] 02:26:00, remote-pe 100.100.100.1, routing-instance
 vpws3006, route-type Egress, vlan-id 400600
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
```

 Copyright © 2018, Juniper Networks, Inc.

```
302720             *[EVPN/7] 02:25:59, remote-pe 100.100.100.3, routing-instance
 vpws3006, route-type Egress, vlan-id 400600
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302736             *[EVPN/7] 02:25:59, remote-pe 100.100.100.3, routing-instance
 vpws3006, route-type Egress Protection, vlan-id 300600
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
ge-0/0/1.1001      *[EVPN/7] 02:29:23
                    > via ge-0/0/1.2001
ge-0/0/1.2001      *[EVPN/7] 02:29:23
                    > via ge-0/0/1.1001
ge-0/0/1.1002      *[EVPN/7] 02:27:06
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
ae10.2003          *[EVPN/7] 02:29:10
                    > via ge-0/0/1.1003
ge-0/0/1.1003      *[EVPN/7] 02:27:06
                    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3

                    > via ae10.2003
ge-0/0/1.1004      *[EVPN/7] 02:27:06
                    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
ae10.1010          *[EVPN/7] 02:27:06
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
ge-0/0/1.3001      *[EVPN/7] 02:26:20
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

                    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
ge-0/0/1.3016      *[EVPN/7] 02:26:13
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
ae10.3011          *[EVPN/7] 02:26:06
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
ae10.3006          *[EVPN/7] 02:25:59
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

                    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
```

show route table mpls.0 protocol ospf

```
user@host> show route table mpls.0 protocol ospf
mpls.0: 29 destinations, 29 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

299952             *[L-OSPF/10] 23:59:42, metric 0
                    > to 10.0.10.70 via lt-1/2/0.14, Pop
                     to 10.0.6.60 via lt-1/2/0.12, Swap 800070, Push 800030(top)
299952(S=0)        *[L-OSPF/10] 23:59:42, metric 0
                    > to 10.0.10.70 via lt-1/2/0.14, Pop
                     to 10.0.6.60 via lt-1/2/0.12, Swap 800070, Push 800030(top)
299968             *[L-OSPF/10] 23:59:48, metric 0
                    > to 10.0.6.60 via lt-1/2/0.12, Pop
```

show route table mpls.0 extensive (PTX Series)

```
user@host> show route table mpls.0 extensive
ge-0/0/2.600 (1 entry, 1 announced)
TSI:
KRT in-kernel ge-0/0/2.600.0      /32 -> {composite(570)}
        *L2VPN  Preference: 7
                Next hop type: Indirect
                Address: 0x9438f34
                Next-hop reference count: 2
                Next hop type: Router, Next hop index: 567
                Next hop: 10.0.0.1 via ge-0/0/1.0, selected
                Label operation: Push 299808
                Label TTL action: prop-ttl
                Load balance label: Label 299808:None;
                Session Id: 0x1
                Protocol next hop: 10.255.255.1
                Label operation: Push 299872 Offset: 252
                Label TTL action: no-prop-ttl
                Load balance label: Label 299872:Flow label PUSH;
                Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
                Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
                State: <Active Int>
                Age: 47         Metric2: 1
                Validation State: unverified
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 2-Common L2 VC
                AS path: I
                Composite next hops: 1
                        Protocol next hop: 10.255.255.1 Metric: 1
                        Label operation: Push 299872 Offset: 252
                        Label TTL action: no-prop-ttl
                        Load balance label: Label 299872:Flow label PUSH;
                        Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
                        Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.0.0.1 via ge-0/0/1.0
                                Session Id: 0x1
                        10.255.255.1/32 Originating RIB: inet.3
                          Metric: 1                       Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: 10.0.0.1 via ge-0/0/1.0
```

show route table mpls.0 (RSVP Route—Transit LSP)

```
user@host> show route table mpls.0

mpls.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 00:37:31, metric 1
                      Receive
1                  *[MPLS/0] 00:37:31, metric 1
                      Receive
2                  *[MPLS/0] 00:37:31, metric 1
                      Receive
13                 *[MPLS/0] 00:37:31, metric 1
                      Receive
300352             *[RSVP/7/1] 00:08:00, metric 1
```


Copyright © 2018, Juniper Networks, Inc.

```
                   > to 10.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300352(S=0)        *[RSVP/7/1] 00:08:00, metric 1
                   > to 10.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300384             *[RSVP/7/2] 00:05:20, metric 1
                    > to 10.64.1.106 via ge-1/0/0.0, Pop
300384(S=0)        *[RSVP/7/2] 00:05:20, metric 1
                    > to 10.64.1.106 via ge-1/0/0.0, Pop
```

show route table vpls_1 detail

```
user@host> show route table vpls_1 detail
vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

172.16.1.11:1000:1:1/96 (1 entry, 1 announced)
*L2VPN Preference: 170/-1
Receive table: vpls_1.l2vpn.0
Next-hop reference count: 2
State: <Active Int Ext>
Age: 4:29:47 Metric2: 1
Task: vpls_1-l2vpn
Announcement bits (1): 1-BGP.0.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:Site-Down
Label-base: 800000, range: 8, status-vector: 0xFF
```

show route table vpn-a

```
user@host> show route table vpn-a
vpn-a.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both
192.168.16.1:1:1:1/96
                   *[VPN/7] 05:48:27
                     Discard
192.168.24.1:1:2:1/96
                   *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
                      AS path: I
                    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
192.168.24.1:1:3:1/96
                   *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
                      AS path: I
                    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

show route table vpn-a.mdt.0

```
user@host> show route table vpn-a.mdt.0
vpn-a.mdt.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:1:0:10.255.14.216:232.1.1.1/144
                   *[MVPN/70] 01:23:05, metric2 1
                      Indirect
1:1:1:10.255.14.218:232.1.1.1/144
                   *[BGP/170] 00:57:49, localpref 100, from 10.255.14.218
                      AS path: I
                    > via so-0/0/0.0, label-switched-path r0e-to-r1
1:1:2:10.255.14.217:232.1.1.1/144
                   *[BGP/170] 00:57:49, localpref 100, from 10.255.14.217
```

```
                      AS path: I
> via so-0/0/1.0, label-switched-path r0-to-r2
```

show route table VPN-A detail

```
user@host> show route table VPN-A detail
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
10.255.179.9/32 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.179.13:200
                Next hop type: Indirect
                Next-hop reference count: 5
                Source: 10.255.179.13
                Next hop type: Router, Next hop index: 732
                Next hop: 10.39.1.14 via fe-0/3/0.0, selected
                Label operation: Push 299824, Push 299824(top)
                Protocol next hop: 10.255.179.13
                Push 299824
                Indirect next hop: 8f275a0 1048574
                State: (Secondary Active Int Ext)
                Local AS: 1 Peer AS: 1
                Age: 3:41:06 Metric: 1 Metric2: 1
                Task: BGP_1.10.255.179.13+64309
                Announcement bits (2): 0-KRT 1-BGP RT Background
                AS path: I
                Communities: target:1:200 rte-type:0.0.0.0:1:0
                Import Accepted
                VPN Label: 299824 TTL Action: vrf-ttl-propagate
                Localpref: 100
                Router ID: 10.255.179.13
                Primary Routing Table bgp.l3vpn.0
```

show route table VPN-AB.inet.0

```
user@host> show route table VPN-AB.inet.0
      VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
      + = Active Route, - = Last Active, * = Both

      10.39.1.0/30       *[OSPF/10] 00:07:24, metric 1
                          > via so-7/3/1.0
      10.39.1.4/30       *[Direct/0] 00:08:42
                          > via so-5/1/0.0
      10.39.1.6/32       *[Local/0] 00:08:46
                           Local
      10.255.71.16/32    *[Static/5] 00:07:24
                          > via so-2/0/0.0
      10.255.71.17/32    *[BGP/170] 00:07:24, MED 1, localpref 100, from
      10.255.71.15
                            AS path: I
                          > via so-2/1/0.0, Push 100020, Push 100011(top)
      10.255.71.18/32    *[BGP/170] 00:07:24, MED 1, localpref 100, from
      10.255.71.15
                            AS path: I
                          > via so-2/1/0.0, Push 100021, Push 100011(top)
      10.255.245.245/32  *[BGP/170] 00:08:35, localpref 100
                            AS path: 2 I
                          > to 10.39.1.5 via so-5/1/0.0
      10.255.245.246/32  *[OSPF/10] 00:07:24, metric 1
                          > via so-7/3/1.0
```


Copyright © 2018, Juniper Networks, Inc.

show route table VPN_blue.mvpn-inet6.0

```
user@host> show route table VPN_blue.mvpn-inet6.0
vpn_blue.mvpn-inet6.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:10.255.2.202:65536:10.255.2.202/432
                   *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
                      AS path: I
                    > via so-0/1/3.0
1:10.255.2.203:65536:10.255.2.203/432
                   *[BGP/170] 00:02:37, localpref 100, from 10.255.2.203
                      AS path: I
                    > via so-0/1/0.0
1:10.255.2.204:65536:10.255.2.204/432
                   *[MVPN/70] 00:57:23, metric2 1
                      Indirect
5:10.255.2.202:65536:128:::192.168.90.2:128:ffff::1/432
                   *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
                      AS path: I
                    > via so-0/1/3.0
6:10.255.2.203:65536:64500:128:::10.12.53.12:128:ffff::1/432
                   *[PIM/105] 00:02:37
                      Multicast (IPv6)
7:10.255.2.202:65536:64500:128:::192.168.90.2:128:ffff::1/432
                   *[MVPN/70] 00:02:37, metric2 1
                      Indirect
```

show route table vrf1.mvpn.0 extensive

```
user@host> show route table vrf1.mvpn.0 extensive
1:10.255.50.77:1:10.255.50.77/240 (1 entry, 1 announced)
        *MVPN   Preference: 70
                PMSI: Flags 0x0: Label 0: RSVP-TE:
Session_13[10.255.50.77:0:25624:10.255.50.77]
                Next hop type: Indirect
                Address: 0xbb2c944
                Next-hop reference count: 360
                Protocol next hop: 10.255.50.77
                Indirect next hop: 0x0 - INH Session ID: 0x0
                State: <Active Int Ext>
                Age: 53:03      Metric2: 1
                Validation State: unverified
                Task: mvpn global task
                Announcement bits (3): 0-PIM.vrf1 1-mvpn global task 2-rt-export

                AS path: I
```

show route table inetflow detail

```
user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Next-hop reference count: 2
                State: <Active Ext>
                Local AS: 64502 Peer AS: 64500
                Age: 4
                Task: BGP_64500.10.12.99.5+3792
                Announcement bits (1): 0-Flow
```

Copyright © 2018, Juniper Networks, Inc.

```
                AS path: 64500 I
                Communities: traffic-rate:0:0
                Validation state: Accept, Originator: 10.12.99.5
                Via: 10.12.44.0/24, Active
                Localpref: 100
                Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
        *Flow   Preference: 5
                Next-hop reference count: 2
                State: <Active>
                Local AS: 64502
                Age: 6:30
                Task: RT Flow
                Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
                AS path: I
                Communities: 1:1
```

```
user@host> show route table green.l2vpn.0 (VPLS Multihoming with FEC 129)
green.l2vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.2:100:10.1.1.2/96 AD
                   *[VPLS/170] 1d 03:11:03, metric2 1
                         Indirect
10.1.1.4:100:10.1.1.4/96 AD
                   *[BGP/170] 1d 03:11:02, localpref 100, from 10.1.1.4
                      AS path: I, validation-state: unverified
                    >    via ge-1/2/1.5
10.1.1.2:100:1:0/96 MH
                   *[VPLS/170] 1d 03:11:03, metric2 1
                         Indirect
10.1.1.4:100:1:0/96 MH
                   *[BGP/170] 1d 03:11:02, localpref 100, from 10.1.1.4
                      AS path: I, validation-state: unverified
                    >    via ge-1/2/1.5
10.1.1.4:NoCtrlWord:5:100:100:10.1.1.2:10.1.1.4/176
                   *[VPLS/7] 1d 03:11:02, metric2 1
                    >    via ge-1/2/1.5
10.1.1.4:NoCtrlWord:5:100:100:10.1.1.4:10.1.1.2/176
                   *[LDP/9] 1d 03:11:02
                         Discard
```

```
user@host> show route table red extensive
red.inet.0: 364481 destinations, 714087 routes (364480 active, 48448 holddown, 1
 hidden)
10.0.0.0/32 (3 entries, 1 announced)
        State: <OnList CalcForwarding>
TSI:
KRT in-kernel 10.0.0.0/32 -> {composite(1048575)} Page 0 idx 1 Type 1 val 0x934342c

    Nexthop: Self
    AS path: [2] I
    Communities: target:2:1
Path 10.0.0.0 from 10.3.0.0 Vector len 4.  Val: 1
        @BGP    Preference: 170/-1
                Route Distinguisher: 2:1
                Next hop type: Indirect
                Address: 0x258059e4
                Next-hop reference count: 2
```

 Copyright © 2018, Juniper Networks, Inc.

```
                Source: 2.2.0.0
                Next hop type: Router
                Next hop: 10.1.1.1 via ge-1/1/9.0, selected
                Label operation: Push 707633
                Label TTL action: prop-ttl
                Session Id: 0x17d8
                Protocol next hop: 10.2.0.0
                Push 16
                Composite next hop: 0x25805988 - INH Session ID: 0x193c
                Indirect next hop: 0x23eea900 - INH Session ID: 0x193c
                State: <Secondary Active Int Ext ProtectionPath ProtectionCand>
                Local AS:     2 Peer AS:     2
                Age: 23         Metric2: 35
                Validation State: unverified
                Task: BGP_172.16.2.0.0+34549
                AS path: I
                Communities: target:2:1
                Import Accepted
                VPN Label: 16
                Localpref: 0
                Router ID: 10.2.0.0
                Primary Routing Table bgp.l3vpn.0
                Composite next hops: 1
                        Protocol next hop: 10.2.0.0 Metric: 35
                        Push 16
                        Composite next hop: 0x25805988 - INH Session ID: 0x193c
                        Indirect next hop: 0x23eea900 - INH Session ID: 0x193c
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.1.1.1 via ge-1/1/9.0
                                Session Id: 0x17d8
                        2.2.0.0/32 Originating RIB: inet.3
                          Metric: 35                      Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: 10.1.1.1 via ge-1/1/9.0
         BGP    Preference: 170/-1
                Route Distinguisher: 2:1
                Next hop type: Indirect
                Address: 0x9347028
                Next-hop reference count: 3
                Source: 10.3.0.0
                Next hop type: Router, Next hop index: 702
                Next hop: 10.1.4.2 via ge-1/0/0.0, selected
                Label operation: Push 634278
                Label TTL action: prop-ttl
                Session Id: 0x17d9
                Protocol next hop: 10.3.0.0
                Push 16
                Composite next hop: 0x93463a0 1048575 INH Session ID: 0x17da
                Indirect next hop: 0x91e8800 1048574 INH Session ID: 0x17da
                State: <Secondary NotBest Int Ext ProtectionPath ProtectionCand>

                Inactive reason: Not Best in its group - IGP metric
                Local AS:     2 Peer AS:     2
                Age: 3:34       Metric2: 70
                Validation State: unverified
                Task: BGP_172.16.3.0.0+32805
                Announcement bits (2): 0-KRT 1-BGP_RT_Background
                AS path: I
                Communities: target:2:1
                Import Accepted
```

```
                VPN Label: 16
                Localpref: 0
                Router ID: 10.3.0.0
                Primary Routing Table bgp.l3vpn.0
                Composite next hops: 1
                        Protocol next hop: 10.3.0.0 Metric: 70
                        Push 16
                        Composite next hop: 0x93463a0 1048575 INH Session ID:
0x17da
                        Indirect next hop: 0x91e8800 1048574 INH Session ID:
0x17da
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.1.4.2 via ge-1/0/0.0
                                Session Id: 0x17d9
                        10.3.0.0/32 Originating RIB: inet.3
                          Metric: 70                      Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: 10.1.4.2 via ge-1/0/0.0
        #Multipath Preference: 255
                Next hop type: Indirect
                Address: 0x24afca30
                Next-hop reference count: 1
                Next hop type: Router
                Next hop: 10.1.1.1 via ge-1/1/9.0, selected
                Label operation: Push 707633
                Label TTL action: prop-ttl
                Session Id: 0x17d8
                Next hop type: Router, Next hop index: 702
                Next hop: 10.1.4.2 via ge-1/0/0.0
                Label operation: Push 634278
                Label TTL action: prop-ttl
                Session Id: 0x17d9
                Protocol next hop: 10.2.0.0
                Push 16
                Composite next hop: 0x25805988 - INH Session ID: 0x193c
               Indirect next hop: 0x23eea900 - INH Session ID: 0x193c Weight 0x1

                Protocol next hop: 10.3.0.0
                Push 16
                Composite next hop: 0x93463a0 1048575 INH Session ID: 0x17da
              Indirect next hop: 0x91e8800 1048574 INH Session ID: 0x17da Weight
 0x4000
                State: <ForwardingOnly Int Ext>
                Inactive reason: Forwarding use only
                Age: 23         Metric2: 35
                Validation State: unverified
                Task: RT
                AS path: I
                Communities: target:2:1
```

show route table bgp.evpn.0 extensive |no-more (EVPN)

```
show route table bgp.evpn.0 extensive | no-more
bgp.evpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
2:1000:10::100::00:aa:aa:aa:aa:aa/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
```

 Copyright © 2018, Juniper Networks, Inc.

```
                Source: 10.2.3.4
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
        State: Local AS: 17  Peer AS:17  Age:21:12  Metric2:1  Validation State:
 unverified
                Task: BGP_17.1.2.3.4+50756
                AS path: I
                Communities: target:1111:8388708 encapsulation0:0:0:0:3
                Import Accepted
                Route Label: 100
                ESI: 00:00:00:00:00:00:00:00:00:00
                Localpref: 100
                Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        1.2.3.4/32 Originating RIB: inet.0
                          Metric: 1                       Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0

2:1000:10::200::00:bb:bb:bb:bb:bb/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
                Source: 10.2.3.4
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
          State: Local AS:17  Peer AS:17  Age:19:43  Metric2:1  Validation
State:unverified
                Task: BGP_17.1.2.3.4+50756
                AS path: I
                Communities: target:2222:22 encapsulation0:0:0:0:3
                Import Accepted
                Route Label: 200
                ESI: 00:00:00:00:00:00:00:00:00:00
                Localpref: 100
                Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        10.2.3.4/32 Originating RIB: inet.0
                          Metric: 1                       Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0

2:1000:10::300::00:cc:cc:cc:cc:cc/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
```

Copyright © 2018, Juniper Networks, Inc.

```
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
                Source: 10.2.3.4
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
          State: Local AS:17  Peer AS:17  Age:17:21  Metric2:1  Validation State:
 unverified  Task: BGP 17,1,2,3,4+50756
          AS path: I
                Communities: target:3333:33 encapsulation0:0:0:0:3
                Import Accepted
                Route Label: 300
                ESI: 00:00:00:00:00:00:00:00:00:00
                Localpref: 100
                Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        10.2.3.4/32 Originating RIB: inet.0
                          Metric: 1                       Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0

3:1000:10::100::1.2.3.4/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
                PMSI: Flags 0x0: Label 100: Type INGRESS-REPLICATION 1.2.3.4
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
                Source: 10.2.3.4
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
          State: Local AS:17  Peer AS:17 Age:37:01  Metric2:1  Validation State:
 unverified  Task: BGP 17.1.2.3.4+50756
                AS path: I
                Communities: target:1111:8388708 encapsulation0:0:0:0:3
                Import Accepted
                Localpref: 100
                Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        10.2.3.4/32 Originating RIB: inet.0
                          Metric: 1                       Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0

3:1000:10::200::1.2.3.4/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
```

 Copyright © 2018, Juniper Networks, Inc.

```
                PMSI: Flags 0x0: Label 200: Type INGRESS-REPLICATION 1.2.3.4
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
                Source: 10.2.3.4
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
          State: Local AS: 17 Peer AS: 17   Age:35:22  Metric2:1  Validation
State:unverified  Task: BGP 17.1.2.3.4+50756
           AS path:I  Communities: target:2222:22 encapsulation):0:0:0:0:3

Import Accepted
                Localpref: 100
                Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        10.2.3.4/32 Originating RIB: inet.0
                          Metric: 1                       Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0

3:1000:10::300::1.2.3.4/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
                PMSI: Flags 0x0: Label 300: Type INGRESS-REPLICATION 1.2.3.4
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
                Source: 10.2.3.4
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
         State: Local AS: 17 Peer AS:  17  Age 35:22 Metric2:1  Validation State:
 unverified   Task: BGP 17.1.2.3.4+5075
          6  AS path: I   Communities: target:3333:33  encapsulation0:0:0:0:3
Import Accepted   Localpref:100
          Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        10.2.3.4/32 Originating RIB: inet.0
                          Metric: 1                       Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0
```

show route terse

List of Syntax

Syntax on page 2416
Syntax (EX Series Switches) on page 2416

Syntax

show route terse
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show route terse

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display a high-level summary of the routes in the routing table.

NOTE: For BGP routes, the **show route terse** command displays the local preference attribute and MED instead of the metric1 and metric2 values. This is mostly due to historical reasons.

To display the metric1 and metric2 value of a BGP route, use the **show route extensive** command.

Options

none—Display a high-level summary of the routes in the routing table.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show route terse on page 2418

Output Fields

Table 216 on page 2416 describes the output fields for the **show route terse** command. Output fields are listed in the approximate order in which they appear.

Table 216: show route terse Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).
number destinations	Number of destinations for which there are routes in the routing table.

 Copyright © 2018, Juniper Networks, Inc.

Table 216: show route terse Output Fields (continued)

Field Name	Field Description
number routes	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active) • **holddown** (routes that are in the pending state before being declared inactive) • **hidden** (routes that are not used because of a routing policy)
route key	Key for the state of the route: • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route.
A	Active route. An asterisk (*) indicates this is the active route.
V	Validation status of the route: • **?**—Not evaluated. Indicates that the route was not learned through BGP. • **I**—Invalid. Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • **N**—Unknown. Indicates that the prefix is not among the prefixes or prefix ranges in the database. • **V**—Valid. Indicates that the prefix and autonomous system pair are found in the database.
Destination	Destination of the route.
P	Protocol through which the route was learned: • **A**—Aggregate • **B**—BGP • **C**—CCC • **D**—Direct • **G**—GMPLS • **I**—IS-IS • **L**—L2CKT, L2VPN, LDP, Local • **K**—Kernel • **M**—MPLS, MSDP • **O**—OSPF • **P**—PIM • **R**—RIP, RIPng • **S**—Static • **T**—Tunnel
Prf	Preference value of the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.

Copyright © 2018, Juniper Networks, Inc.

Table 216: show route terse Output Fields (continued)

Field Name	Field Description
Metric 1	First metric value in the route. For routes learned from BGP, this is the MED metric.
Metric 2	Second metric value in the route. For routes learned from BGP, this is the IGP metric.
Next hop	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated.

Sample Output

show route terse

```
user@host> show route terse
inet.0: 10 destinations, 12 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A V Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* ? 172.16.1.1/32        O  10          1             >10.0.0.2
  ?                    B 170        100                             I
  unverified                                       >10.0.0.2
* ? 172.16.1.1/32        D   0                        >lo0.2
* V 2.2.0.2/32         B 170        110                             200 I
  valid                                            >10.0.0.2
* ? 10.0.0.0/30        D   0                       >lt-1/2/0.1
  ?                    B 170        100                             I
  unverified                                       >10.0.0.2
* ? 10.0.0.1/32        L   0                        Local
* ? 10.0.0.4/30        B 170        100                             I
  unverified                                       >10.0.0.2
* ? 10.0.0.8/30        B 170        100                             I
  unverified                                       >10.0.0.2
* I 172.16.1.1/32      B 170         90                             200 I
  invalid                                          >10.0.0.2
* N 192.168.2.3/32     B 170        100                             200 I
  unknown                                          >10.0.0.2
* ? 172.16.233.5/32      O  10          1              MultiRecv
```

 Copyright © 2018, Juniper Networks, Inc.

CHAPTER 26

RIP Operational Commands

clear rip general-statistics

List of Syntax
Syntax on page 2420
Syntax (EX Series Switches and QFX Series) on page 2420

Syntax

clear rip general-statistics
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches and QFX Series)

clear rip general-statistics

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Clear RIP general statistics.

Options

none—Clear RIP general statistics.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Related Documentation

- show rip general-statistics on page 2423

List of Sample Output
clear rip general-statistics on page 2420

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear rip general-statistics

```
user@host> clear rip general-statistics
```

 Copyright © 2018, Juniper Networks, Inc.

clear rip statistics

List of Syntax

Syntax on page 2421
Syntax (EX Series Switches and QFX Series) on page 2421

Syntax

clear rip statistics
<instance (all | *instance-name*)>
<logical-system (all | *logical-system-name*)>
<*neighbor*>
<peer (all | *address*)>

Syntax (EX Series Switches and QFX Series)

clear rip statistics
<instance (all | *instance-name*)>
<*neighbor*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Clear RIP statistics.

Options

none—Reset RIP counters for all neighbors for all routing instances.

instance (all | *instance-name*)—(Optional) Clear RIP statistics for all instances or for the specified routing instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

neighbor—(Optional) Clear RIP statistics for the specified neighbor only.

peer (all | *address*)—(Optional) Clear RIP statistics for a single peer or all peers.

Required Privilege Level

clear

Related Documentation

- show rip statistics on page 2427

List of Sample Output

clear rip statistics on page 2421

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear rip statistics

```
user@host> clear rip statistics
```

Copyright © 2018, Juniper Networks, Inc.

show rip general-statistics

List of Syntax

Syntax on page 2423
Syntax (EX Series Switches and QFX Series) on page 2423

Syntax

show rip general-statistics
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches and QFX Series)

show rip general-statistics

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display brief RIP statistics.

Options

none—Display brief RIP statistics.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- clear rip general-statistics on page 2420

List of Sample Output

show rip general-statistics on page 2424

Output Fields

Table 217 on page 2423 lists the output fields for the **show rip general-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 217: show rip general-statistics Output Fields

Field Name	Field Description
bad msgs	Number of invalid messages received.
no recv intf	Number of packets received with no matching interface.
curr memory	Amount of memory currently used by RIP.
max memory	Most memory used by RIP.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show rip general-statistics

```
user@host> show rip general-statistics
RIPv2 I/O info:
     bad msgs       :         0
     no recv intf   :         0
     curr memory    :         0
     max memory     :         0
```


Copyright © 2018, Juniper Networks, Inc.

show rip neighbor

List of Syntax

Syntax on page 2425
Syntax (EX Series Switches and QFX Series) on page 2425

Syntax

show rip neighbor
<instance (all | *instance-name*)>
<logical-system (all | *logical-system-name*)>
<*name*>

Syntax (EX Series Switches and QFX Series)

show rip neighbor
<instance (all | *instance-name*)>
<*name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display information about RIP neighbors.

Options

none—Display information about all RIP neighbors for all instances.

instance (all | *instance-name*)—(Optional) Display RIP neighbor information for all instances or for only the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

name—(Optional) Display detailed information about only the specified RIP neighbor.

Required Privilege Level

view

List of Sample Output

show rip neighbor on page 2426
show rip neighbor (With Demand Circuits Configured) on page 2426

Output Fields

Table 218 on page 2426 lists the output fields for the **show rip neighbor** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2018, Juniper Networks, Inc.

Table 218: show rip neighbor Output Fields

Field Name	Field Description
Neighbor	Name of the RIP neighbor. NOTE: Beginning with Junos OS Release 11.1, when you configure demand circuits, the output displays a demand circuit (DC) flag next to neighbor interfaces configured for demand circuits. If you configure demand circuits at the **[edit protocols rip group *group-name* neighbor *neighbor-name*]** hierarchy level, the output shows only the neighboring interface that you specifically configured as a demand circuit. If you configure demand circuits at the **[edit protocols rip group *group-name*]** hierarchy level, all of the interfaces in the group are configured as demand circuits. Therefore, the output shows all of the interfaces in that group as demand circuits.
State	State of the connection: **Up** or **Dn** (Down).
Source Address	Address of the port on the local router.
Destination Address	Address of the port on the remote router.
Send Mode	Send options: **broadcast**, **multicast**, **none**, or **version 1**.
Receive Mode	Type of packets to accept: **both**, **none**, **version 1**, or **version 2**.
In Met	Metric added to incoming routes when advertising into RIP routes that were learned from other protocols.

Sample Output

show rip neighbor

```
user@host> show rip neighbor
                  Local  Source          Destination     Send   Receive  In
Neighbor          State  Address         Address         Mode   Mode     Met
--------          -----  -------         -----------     ----   -------  ---
ge-2/3/0.0           Up  192.168.9.105   192.168.9.107   bcast  both       1
at-5/1/1.42          Dn  (null)          (null)          mcast  v2 only    3
at-5/1/0.42          Dn  (null)          (null)          mcast  both       3
at-5/1/0.0           Up  198.51.100.0    224.0.0.9       mcast  both       3
so-0/0/0.0           Up  192.168.9.97    224.0.0.9       mcast  both       3
```

show rip neighbor (With Demand Circuits Configured)

```
user@host> show rip neighbor
                  Local  Source          Destination     Send   Receive  In
Neighbor          State  Address         Address         Mode   Mode     Met
--------          -----  -------         -----------     ----   -------  ---
so-0/1/0.0(DC)       Up  10.10.10.2      224.0.0.9       mcast  both       1
so-0/2/0.0(DC)       Up  192.0.2.2       224.0.0.9       mcast  both       1
```


Copyright © 2018, Juniper Networks, Inc.

show rip statistics

List of Syntax

Syntax on page 2427
Syntax (EX Series Switches and QFX Series) on page 2427

Syntax

show rip statistics
<instance (all | *instance-name*)>
<logical-system (all | *logical-system-name*)>
<*name*>
<peer (all | *address*)>

Syntax (EX Series Switches and QFX Series)

show rip statistics
<instance (all | *instance-name*)>
<*name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display RIP statistics about messages sent and received on an interface, as well as information received from advertisements from other routing devices.

Options

none—Display RIP statistics for all routing instances.

instance (all | *instance-name*)—(Optional) Display RIP statistics for all instances or for only the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

name—(Optional) Display detailed information about only the specified RIP neighbor.

peer (all | *address*)—(Optional) Display RIP statistics for a single peer or all peers.

Required Privilege Level

view

Related Documentation

- clear rip statistics on page 2421

List of Sample Output

show rip statistics on page 2429

Output Fields

Table 219 on page 2428 lists the output fields for the **show rip statistics** command. Output fields are listed in the approximate order in which they appear.

Table 219: show rip statistics Output Fields

Field Name	Field Description
RIP info	Information about RIP on the specified interface: • **port**—UDP port number used for RIP. • **update interval**—Interval between routing table updates, in seconds. • **holddown**—Hold-down interval, in seconds. • **timeout**—Timeout interval, in seconds. • **restart in progress**—Graceful restart status. Displayed when RIP is or has been in the process of graceful restart. • **restart time**—Estimated time for the graceful restart to finish, in seconds. • **restart will complete in**—Remaining time for the graceful restart to finish, in seconds. • **rts learned**—Number of routes learned through RIP. • **rts held down**—Number of routes held down by RIP. • **rqsts dropped**—Number of received request packets that were dropped. • **resps dropped**—Number of received response packets that were dropped.
logical-interface	Name of the logical interface and its statistics: • **routes learned**—Number of routes learned on the logical interface. • **routes advertised**—Number of routes advertised by the logical interface.
Counter	List of counter types: • **Updates Sent**—Number of update messages sent. • **Triggered Updates Sent**—Number of triggered update messages sent. • **Responses Sent**—Number of response messages sent. • **Bad Messages**—Number of invalid messages received. • **RIPv1 Updates Received**—Number of RIPv1 update messages received. • **RIPv1 Bad Route Entries**—Number of RIPv1 invalid route entry messages received. • **RIPv1 Updates Ignored**—Number of RIPv1 update messages ignored. • **RIPv2 Updates Received**—Number of RIPv2 update messages received. • **RIPv2 Bad Route Entries**—Number of RIPv2 invalid route entry messages received. • **RIPv2 Updates Ignored**—Number of RIPv2 update messages that were ignored. • **Authentication Failures**—Number of received update messages that failed authentication. • **RIP Requests Received**—Number of RIP request messages received. • **RIP Requests Ignored**—Number of RIP request messages ignored.
Total	Total number of packets for the selected counter.
Last 5 min	Number of packets for the selected counter in the most recent 5-minute period.
Last minute	Number of packets for the selected counter in the most recent 1-minute period.

 Copyright © 2018, Juniper Networks, Inc.

Sample Output

show rip statistics

```
user@host> show rip statistics so-0/0/0.0
RIP info: port 520; update interval: 30s; holddown 180s; timeout 120s
restart in progress: restart time 60s; restart will complete in 55s
    rts learned  rts held down  rqsts dropped  resps dropped
              0              0              0              0
so-0/0/0.0:  0 routes learned; 501 routes advertised
Counter                         Total   Last 5 min  Last minute
-------                   -----------  -----------  -----------
Updates Sent                        0            0            0
Triggered Updates Sent              0            0            0
Responses Sent                      0            0            0
Bad Messages                        0            0            0
RIPv1 Updates Received              0            0            0
RIPv1 Bad Route Entries             0            0            0
RIPv1 Updates Ignored               0            0            0
RIPv2 Updates Received              0            0            0
RIPv2 Bad Route Entries             0            0            0
RIPv2 Updates Ignored               0            0            0
Authentication Failures             0            0            0
RIP Requests Received               0            0            0
RIP Requests Ignored                0            0            0
```

Copyright © 2018, Juniper Networks, Inc.

 Copyright © 2018, Juniper Networks, Inc.

CHAPTER 27

RIPng Operational Commands

clear ripng general-statistics

List of Syntax

Syntax on page 2432
Syntax (EX Series Switches) on page 2432

Syntax

clear ripng general-statistics
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

clear ripng general-statistics

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Clear RIP next generation (RIPng) general statistics.

Options

none—Clear RIPng general statistics.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Related Documentation

- show ripng general-statistics on page 2434

List of Sample Output

clear ripng general-statistics on page 2432

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ripng general-statistics

```
user@host> clear ripng general-statistics
```

 Copyright © 2018, Juniper Networks, Inc.

clear ripng statistics

List of Syntax

Syntax on page 2433
Syntax (EX Series Switch) on page 2433

Syntax

clear ripng statistics
<*instance* | *name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch)

clear ripng statistics
<*instance* | *name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Clear RIP next-generation (RIPng) statistics.

Options

none—Reset RIPng counters for all neighbors for all routing instances.

instance—(Optional) Reset RIPng counters for the specified instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

name—(Optional) Reset RIPng counters for the specified neighbor.

Required Privilege Level

clear

Related Documentation

- show ripng statistics on page 2438

List of Sample Output

clear ripng statistics on page 2433

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ripng statistics

```
user@host> clear ripng statistics
```

show ripng general-statistics

List of Syntax

Syntax on page 2434
Syntax (EX Series Switch) on page 2434

Syntax

show ripng general-statistics
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch)

show ripng general-statistics

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display general RIP next-generation (RIPng) statistics.

Options

none—Display general RIPng statistics.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- clear ripng general-statistics on page 2432

List of Sample Output

show ripng general-statistics on page 2435

Output Fields

Table 220 on page 2434 lists the output fields for the **show ripng general-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 220: show ripng general-statistics Output Fields

Field Name	Field Description
bad msgs	Number of invalid messages received.
no recv intf	Number of packets received with no matching interface.
curr memory	Amount of memory currently used by RIPng.
max memory	Most memory used by RIPng.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show ripng general-statistics

```
user@host> show ripng general-statistics
RIPng I/O info:
     bad msgs       :         0
     no recv intf   :         0
     curr memory    :         0
     max memory     :         0
```

show ripng neighbor

List of Syntax

Syntax on page 2436
Syntax (EX Series Switch) on page 2436

Syntax

show ripng neighbor
<logical-system (all | *logical-system-name*)>
<*name*>

Syntax (EX Series Switch)

show ripng neighbor
<*name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display information about RIP next-generation (RIPng) neighbors.

Options

none—Display information about all RIPng neighbors.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

name—(Optional) Display detailed information about a specific RIPng neighbor.

Required Privilege Level

view

List of Sample Output

show ripng neighbor on page 2437

Output Fields

Table 221 on page 2436 lists the output fields for the **show ripng neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 221: show ripng neighbor Output Fields

Field Name	Field Description
Neighbor	Name of RIPng neighbor.
State	State of the connection: **Up** or **Dn** (Down).
Source Address	Source address.
Destination Address	Destination address.
Send	Send options: **broadcast**, **multicast**, **none**, **version 1**, or **yes.**
Recv	Type of packets to accept: **both**, **none**, **version 1**, or **yes.**

 Copyright © 2018, Juniper Networks, Inc.

Table 221: show ripng neighbor Output Fields (continued)

Field Name	Field Description
In Met	Metric added to incoming routes when advertising into RIPng routes that were learned from other protocols.

Sample Output

show ripng neighbor

```
user@host> show ripng neighbor
                        Source                                 Dest                 In
Neighbor        State   Address                                Address   Send Recv  Met
--------        -----   -------                                -------   ---- ----  ---
fe-0/0/2.0         Up   fe80::290:69ff:fe68:b002               ff02::9   yes  yes     1
```

Copyright © 2018, Juniper Networks, Inc.

show ripng statistics

List of Syntax

Syntax on page 2438
Syntax (EX Series Switch) on page 2438

Syntax

show ripng statistics
<logical-system (all | *logical-system-name*)>
<*name*>

Syntax (EX Series Switch)

show ripng statistics
<*name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display RIP next generation (RIPng) statistics about messages sent and received on an interface, as well as information received from advertisements from other routing devices.

Options

none—Display RIPng statistics for all neighbors.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

name—(Optional) Display detailed information about a specific RIPng neighbor.

Required Privilege Level

view

Related Documentation

- clear ripng statistics on page 2433

List of Sample Output

show ripng statistics on page 2439

Output Fields

Table 222 on page 2439 lists the output fields for the **show ripng statistics** command. Output fields are listed in the approximate order in which they appear.

 Copyright © 2018, Juniper Networks, Inc.

Table 222: show ripng statistics Output Fields

Field Name	Field Description
RIPng info	Information about RIPng on the specified interface: • **port**—UDP port number used for RIPng. • **holddown**—Hold-down interval, in seconds. • **rts learned**—Number of routes learned through RIPng. • **rts held down**—Number of routes held down by RIPng. • **rqsts dropped**—Number of received request packets that were dropped. • **resps dropped**—Number of received response packets that were dropped. • **restart**—Graceful restart status. Displayed when RIPng is or has been in the process of graceful restart.
logical-interface	Name of the logical interface and its statistics: • **routes learned**—Number of routes learned on the logical interface. • **routes advertised**—Number of routes advertised by the logical interface. • **timeout**—Timeout interval, in seconds. • **update interval**—Interval between routing table updates, in seconds.
Counter	List of counter types: • **Updates Sent**—Number of update messages sent. • **Triggered Updates Sent**—Number of triggered update messages sent. • **Responses Sent**—Number of response messages sent. • **Bad Messages**—Number of invalid messages received. • **Updates Received**—Number of RIPng update messages received. • **Bad Route Entries**—Number of RIPng invalid route entry messages received. • **Updates Ignored**—Number of RIPng update messages ignored. • **RIPng Requests Received**—Number of RIPng request messages received. • **RIPng Requests Ignored**—Number of RIPng request messages ignored.
Total	Total number of packets for the selected counter.
Last 5 min	Number of packets for the selected counter in the most recent 5-minute period.
Last minute	Number of packets for the selected counter in the most recent 1-minute period.

Sample Output

show ripng statistics

```
user@host> show ripng statistics
RIPng info: port 521; holddown 120s;
    rts learned  rts held down  rqsts dropped  resps dropped
              0              0              0              0

so-0/1/3.0:  0 routes learned; 1 routes advertised; timeout 180s; update interval
 20s
Counter                          Total   Last 5 min  Last minute
-------                    -----------  -----------  -----------
```

```
Updates Sent                    934         16          4
Triggered Updates Sent            1          0          0
Responses Sent                    0          0          0
Bad Messages                      0          0          0
Updates Received                  0          0          0
Bad Route Entries                 0          0          0
Updates Ignored                   0          0          0
RIPng Requests Received           0          0          0
RIPng Requests Ignored            0          0          0
```

 Copyright © 2018, Juniper Networks, Inc.

CHAPTER 28

Firewall Filter Operational Commands

Copyright © 2018, Juniper Networks, Inc.

clear firewall

List of Syntax

Syntax on page 2442
Syntax (EX Series Switches) on page 2442

Syntax

clear firewall (all | counter *counter-name* | filter *filter-name* | log (all | *logical-system-name*) | logical-system *logical-system-name*)

Syntax (EX Series Switches)

clear firewall (all | counter *counter-name* | filter *filter-name* | log (all | *logical-system-name*) | policer counter (all | counter-id *counter-index*))

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
logical-system option introduced in Junos OS Release 9.3.
log option introduced before Junos OS Release 11.4.

Description

Clear statistics about configured firewall filters.

When you clear the counters of a filter, this impacts not only the counters shown by the CLI, but also the ones tracked by SNMP2.

Subscriber management uses firewall filters to capture and report the volume-based service accounting counters that are used for subscriber billing. The **clear firewall** command also clears the service accounting counters that are reported to the RADIUS accounting server. For this reason, you must be cautious in specifying which firewall statistics you want to clear.

NOTE: The clear firewall command cannot be used to clear the Routing Engine filter counters on a backup Routing Engine that is enabled for graceful Routing Engine switchover (GRES).

If you clear statistics for firewall filters that are applied to Trio-based DPCs and that also use the **prefix-action** action on matched packets, wait at least 5 seconds before you enter the **show firewall prefix-action-stats** command. A 5-second pause between issuing the **clear firewall** and **show firewall prefix-action-stats** commands avoids a possible timeout of the **show firewall prefix-action-stats** command.

Options

all—Clear the packet and byte counts for all filters. On EX Series switches, this option also clears the packet counts for all policer counters.

counter *counter-name*—Clear the packet and byte counts for a filter counter that has been configured with the counter firewall filter action.

filter *filter-name*—Clear the packet and byte counts for the specified firewall filter.

 Copyright © 2018, Juniper Networks, Inc.

log (all | *logical-system-name*)—Clear log entries for IPv4 firewall filters that have **then log** as an action. Use **log all** to clear all log entries or **log *logical-system-name*** to clear log entries for the specified logical system.

logical-system *logical-system-name*—Clear the packet and byte counts for the specified logical system.

policer counter (all | counter-id *counter-index*)—(EX8200 switches only) Clear all policer counters using the **policer counter all** command, or clear a specific policer counter using the **policer counter counter-id *counter-index*** command. The value of *counter-index* can be 0, 1, or 2.

Required Privilege Level clear

Related Documentation

- show firewall on page 2444

List of Sample Output

Sample Output

clear firewall all

```
user@host> clear firewall all
```

clear firewall (counter counter-name)

```
user@host> clear firewall counter port-filter-counter
```

clear firewall (filter filter-name)

```
user@host> clear firewall filter ingress-port-filter
```

clear firewall (policer counter all) (EX8200 Switch)

```
user@switch> clear firewall policer counter all
```

clear firewall (policer counter counter-id counter-index) (EX8200 Switch)

```
user@switch> clear firewall policer counter counter-id 0
```

Copyright © 2018, Juniper Networks, Inc.

show firewall

List of Syntax

Syntax on page 2444
Syntax (EX Series Switches) on page 2444

Syntax

show firewall
<application (CFM | eswd | RMPS)>>
<counter *counter-name*>
<detail>
<filter *filter-name*>
<filter regex *regular-expression*>
<logical-system (all | *logical-system-name*)>
<terse>

Syntax (EX Series Switches)

show firewall
<application (CFM | eswd | RMPS)>>
<counter *counter-name*>
<detail>
<filter *filter-name*>
<filter regex *regular-expression*>
<log <(detail | interface *interface-name*)>>
<policer counters <(detail | counter-id *counter-index* <detail>)>>
<terse>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Option **logical-system** introduced in Junos OS Release 9.3.
Option **terse** introduced in Junos OS Release 9.4.
Option **policer counters** introduced in Junos OS Release 12.2 for EX Series switches.
Option **detail** introduced in Junos OS Release 12.3 for EX Series switches.
Option **detail** introduced in Junos OS Release 14.1 for MX Series routers.
Option **regex *regular-expression*** introduced in Junos OS Release 14.2.

Description

Display enhanced statistics and counters for all configured firewall filters.

Options

none—(Optional) Display statistics and counters for all configured firewall filters and counters. For EX Series switches, this command also displays statistics about all configured policers.

application (CFM | eswd | RMPS)—(Optional) Show firewall elements owned by the selected software component:

- Connectivity Fault Management (CFM)
- Ethernet switching daemon (eswd)—Shows only on devices that support it.
- Resource Management and Packet Steering (RMPS)

counter *counter-name*—(Optional) Name of a filter counter.

Copyright © 2018, Juniper Networks, Inc.

detail—(EX Series switches and MX Series routers only) (Optional) Display firewall filter statistics and enhanced policer statistics and counters.

filter *filter-name*—(Optional) Name of a configured filter.

filter regex *regular-expression*—(Optional) Regular expression that matches the names of a subset of filters.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

log—(Optional) Display log entries for firewall filters.

log <(detail | interface *interface-name*)>—(EX Series switches only) (Optional) Display detailed log entries of firewall activity or log information about a specific interface.

policer counters <(detail | counter-id *counter-index* <detail>)>—(EX8200 switches only) (Optional) Display enhanced policer counter statistics in brief or in detail.

terse—(Optional) Display firewall filter names only.

Required Privilege Level view

Related Documentation

- clear firewall on page 2442
- show firewall log on page 2453
- *Verifying That Firewall Filters Are Operational*
- *Verifying That Policers Are Operational*
- show policer on page 2460
- *Enhanced Policer Statistics Overview*
- *enhanced-policer*

List of Sample Output

show firewall filter (MX Series Router and EX Series Switch) on page 2448
show firewall filter (non MX Series Router and EX Series Switch) on page 2448
show firewall filter (Dynamic Input Filter) on page 2448
show firewall (Logical Systems) on page 2448
show firewall (counter counter-name) on page 2449
show firewall log on page 2449
show firewall policer counters (EX8200 Switch) on page 2449
show firewall policer counters (detail) (EX8200 Switch) on page 2449
show firewall policer counters (counter-id counter-index) (EX8200 Switch) on page 2450
show firewall policer counters (counter-id counter-index detail) (EX8200 Switch) on page 2450
show firewall detail on page 2450

Output Fields Table 223 on page 2446 lists the output fields for the **show firewall** command. Output fields are listed in the approximate order in which they appear.

Table 223: show firewall Output Fields

Field Name	Field Description
Filter	Name of a filter that has been configured with the **filter** statement at the **[edit firewall]** hierarchy level. Except on EX Series switches: • When an interface-specific filter is displayed, the name of the filter is followed by the full interface name and by either **-i** for an input filter or **-o** for an output filter. • When dynamic filters are displayed, the name of the filter is followed by the full interface name and by either **-in** for an input filter or **-out** for an output filter. When a logical system–specific filter is displayed, the name of the filter is prefixed with two underscore (__) characters and the name of the logical system (for example, **__ls1/filter1**). • When a service filter is displayed that uses a service set, the separator between the service-set name and the service-filter name is a semicolon (:). NOTE: For **bridge family filter**, the **ip-protocol** match criteria is supported only for IPv4 and not for IPv6. This is applicable for line cards that support the Junos Trio chipset, such as the MX 3D MPC line cards.
Counters	Display filter counter information: • **Name**—Name of a filter counter that has been configured with the **counter** firewall filter action. • **Bytes**—Number of bytes that match the filter term under which the **counter** action is specified. • **Packets**—Number of packets that matched the filter term under which the **counter** action is specified. NOTE: On M and T Series routers, firewall filters cannot count **ip-options** packets on a per option type and per interface basis. A limited work around is to use the **show pfe statistics ip options** command to see **ip-options** statistics on a per Packet Forwarding Engine (PFE) basis. See *show pfe statistics ip* for sample output.
Policers	Display policer information: • **Name**—Name of policer. • **Bytes**—(For two-color policers on MX Series routers and EX Series switches, and for hierarchical policers on interfaces hosted on MICs and MPCs in MX Series routers) Number of bytes that match the filter term under which the policer action is specified. This is only the number out-of-specification (out-of-spec) byte counts, not all the bytes in all packets policed by the policer. For other combinations of policer type, device, and line card type, this field is blank. • **Packets**—Number of packets that matched the filter term under which the policer action is specified. This is only the number of out-of-specification (out-of-spec) packet counts, not all packets policed by the policer.
Policer Counter Index	(EX8200 switch only) Global management counter ID. The counter ID value (*counter-index*) can be 0, 1, or 2.
Green	(EX8200 switch only) Number of packets within the limits. The number of packets is smaller than the committed information rate (CIR).
Yellow	(EX8200 switch only) Number of packets partially within the limits. The number of packets is greater than the CIR, but the burst size is within the excess burst size (EBS) limit.

Copyright © 2018, Juniper Networks, Inc.

Table 223: show firewall Output Fields (continued)

Field Name	Field Description
Discard	(EX8200 switch only) Number of discarded packets.
Bytes	(EX8200 switch only) Number of green, yellow, red, or discarded packets in bytes.
Packets	(EX8200 switch only) Number of green, yellow, red, or discarded packets.
Filter name	(EX8200 switch only) Name of the filter with a term associated to a policer.
Term name	(EX8200 switch only) Name of the term associated with a policer.
Policer name	(EX8200 switch only) Name of the policer that is associated with a global management counter.
P1-t1	• OOS packet statistics for packets that are marked out-of-specification (out-of-spec) by the policer. Changes to all packets that have out-of-spec actions, such as discard, color marking, or forwarding-class, are included in this counter. • Offered packet statistics for traffic subjected to policing. • Transmitted packet statistics for traffic that is not discarded by the policer. When the policer action is discard, the statistics are the same as the in-spec statistics; when the policer action is non-discard (loss-priority or forwarding-class), the statistics are included in this counter.

Sample Output

show firewall filter (MX Series Router and EX Series Switch)

```
user@host> show firewall filter test
Filter: test
Counters:
Name                                Bytes                 Packets
Counter-1                               0                       0
Counter-2                               0                       0
Policers:
Name                                Bytes                 Packets
Policer-1                            2770                      70
```

show firewall filter (non MX Series Router and EX Series Switch)

```
user@host> show firewall filter test
Filter: test
Counters:
Name                                Bytes                 Packets
Counter-1                               0                       0
Counter-2                               0                       0
Policers:
Name                                Bytes                 Packets
Policer-1                                                      70
```

show firewall filter (Dynamic Input Filter)

```
user@host> show firewall filter dfwd-ge-5/0/0.1-in
Filter: dfwd-ge-5/0/0.1-in
Counters:
Name                                                 Bytes                Packets
c1-ge-5/0/0.1-in                                         0                      0
```

show firewall (Logical Systems)

```
user@host> show firewall

Filter: __lr1/test
Counters:
Name                                                 Bytes                Packets
icmp                                                   420                      5
Filter: __default_bpdu_filter__
Filter: __lr1/inet_filter1
Counters:
Name                                                 Bytes                Packets
inet_tcp_count                                           0                      0
inet_udp_count                                           0                      0
Filter: __lr1/inet_filter2
Counters:
Name                                                 Bytes                Packets
inet_icmp_count                                          0                      0
inet_pim_count                                           0                      0
Filter: __lr2/inet_filter1
Counters:
Name                                                 Bytes                Packets
inet_tcp_count                                           0                      0
```

 Copyright © 2018, Juniper Networks, Inc.

```
inet_udp_count                                         0                    0
```

show firewall (counter counter-name)

```
user@host> show firewall counter icmp-counter
Filter: ingress-port-voip-class-filter
Counters:
Name                                                Bytes              Packets
icmp-counter                                            0                    0
```

show firewall log

```
user@host> show firewall log
Log :

Time      Filter    Action Interface     Protocol        Src Addr
          Dest Addr
08:00:53  pfe       R      ge-1/0/1.0    ICMP            192.168.3.5
        192.168.3.4
08:00:52  pfe       R      ge-1/0/1.0    ICMP            192.168.3.5
        192.168.3.4
08:00:51  pfe       R      ge-1/0/1.0    ICMP            192.168.3.5
        192.168.3.4
08:00:50  pfe       R      ge-1/0/1.0    ICMP            192.168.3.5
        192.168.3.4
08:00:49  pfe       R      ge-1/0/1.0    ICMP            192.168.3.5
        192.168.3.4
08:00:48  pfe       R      ge-1/0/1.0    ICMP            192.168.3.5
        192.168.3.4
08:00:47  pfe       R      ge-1/0/1.0    ICMP            192.168.3.5
        192.168.3.4
```

show firewall policer counters (EX8200 Switch)

```
user@switch> show firewall policer counters
Policer Counter Index 0:
                          Bytes            Packets
Green:                       73              15914
Yellow:                       9               1962
Discard:                    119              25942

Policer Counter Index 1:
                           Bytes            Packets
Green:                         0                  0
Yellow:                        0                  0
Discard:                       0                  0

Policer Counter Index 2:
                           Bytes            Packets
Green:                         0                  0
Yellow:                        0                  0
Discard:                       0                  0
```

show firewall policer counters (detail) (EX8200 Switch)

```
user@switch> show firewall policer counters detail
Policer Counter Index 0:
                          Bytes            Packets
```

Copyright © 2018, Juniper Networks, Inc.

```
Green:                     73             15914
Yellow:                     9              1962
Discard:                  119             25942

Filter name                Term name           Policer name
myfilter                   polcr-term-1        myfilter-polcr-1
inet-filter-ae             ae-snmp             policer-1
inet-filter-ae             ae-ssh              policer-2

Policer Counter Index 1:
                          Bytes            Packets
Green:                      0                  0
Yellow:                     0                  0
Discard:                    0                  0

Filter name                Term name           Policer name

Policer Counter Index 2:
                          Bytes            Packets
Green:                      0                  0
Yellow:                     0                  0
Discard:                    0                  0

Filter name                Term name           Policer name
```

show firewall policer counters (counter-id counter-index) (EX8200 Switch)

```
user@switch> show firewall policer counters counter-id 0
Policer Counter Index 0:
                          Bytes            Packets
Green:                     73             15914
Yellow:                     9              1962
Discard:                  119             25942
```

show firewall policer counters (counter-id counter-index detail) (EX8200 Switch)

```
user@switch> show firewall policer counters counter-id 0 detail
Policer Counter Index 0:
                          Bytes            Packets
Green:                     73             15914
Yellow:                     9              1962
Discard:                  119             25942

Filter name                Term name           Policer name
myfilter                   polcr-term-1        myfilter-polcr-1
inet-filter-ae             ae-snmp             policer-1
inet-filter-ae             ae-ssh              policer-2
```

show firewall detail

```
user@host> show firewall detail
Filter: __default_bpdu_filter__

Filter: foo
Counters:
Name                                              Bytes              Packets
c1                                             17652140               160474
Policers:
Name                                              Bytes              Packets
P1-t1
```

 Copyright © 2018, Juniper Networks, Inc.

```
OOS                                                     0                18286
Offered                                                 0 18446744073709376546
Transmitted                                             0 18446744073709358260
```

show firewall filter version

Syntax

show firewall filter version <*filter-name*>

Release Information

Command introduced in Junos OS Release 10.2R2.

Description

Display the version number of the installed firewall filter in the Routing Engine.

Options

none—(Optional) Display the version number of all installed firewall filters.

filter-name—(Optional) Name of a configured filter. If you specify the name of a filter, only the version number of that filter is displayed.

Additional Information

The initial version number is 1. This number increments by one when you modify the firewall filter settings or an associated prefix action. The maximum version number is 4,294,967,295. When the version number reaches 4,294,967,295, this number is reset to 1.

Required Privilege Level

view

List of Sample Output

show firewall filter version on page 2452

Output Fields

Table 224 on page 2452 lists the output fields for the **show firewall filter version** command. Output fields are listed in the approximate order in which they appear.

Table 224: show firewall filter version Output Fields

Field Name	Field Description
Filter	Name of a filter that has been configured with the **filter** statement at the **[edit firewall]** hierarchy level.
Version	Display the version number of the firewall filter.

Sample Output

show firewall filter version

```
user@host> show firewall filter version
Filter version information :
Filter                                              Version
test                                                     10
```


Copyright © 2018, Juniper Networks, Inc.

show firewall log

List of Syntax

Syntax on page 2453
Syntax (EX Series Switches) on page 2453

Syntax

show firewall log
<detail>
<extensive>
<interface *interface-name*>
<logical-system (*logical-system-name* | all)>

Syntax (EX Series Switches)

show firewall log
<detail>
<interface *interface-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
extensive option introduced in Junos OS Release 16.1.
logical-system option introduced in Junos OS Release 9.3.

Description

Display log information about firewall filters.

Options

none—Display log information about firewall filters.

detail—(Optional) Display detailed information.

extensive—(Optional) Display hex dump of packet captured by log action.

interface *interface-name*—(Optional) Display log information about a specific interface.

logical-system (*logical-system-name* | all)—(Optional) Perform this operation on all logical systems or on a particular system.

Required Privilege Level

view

List of Sample Output

show firewall log on page 2454
show firewall log detail on page 2454
show firewall log extensive on page 2455

Output Fields

Table 225 on page 2453 lists the output fields for the **show firewall log** command. Output fields are listed in the approximate order in which they appear.

Table 225: show firewall log Output Fields

Field Name	Field Description
Time of Log	Time that the event occurred.

Copyright © 2018, Juniper Networks, Inc.

Table 225: show firewall log Output Fields (continued)

Field Name	Field Description
Filter	• Displays the name of a configured firewall filter or service filter only if the packet hit the filter's **log** action in a kernel filter (in the control plane). For any traffic that reaches the Routing Engine, the packets hit the **log** action in the kernel. • For all other logged packets (packet hit the filter's **log** action in the Packet Forwarding Engine), this field displays **pfe** instead of a configured filter name.
Filter Action	Filter action: • **A**—Accept • **D**—Discard • **R**—Reject
Name of Interface	• Displays a physical interface name if the packet arrived at a port on a line card. • Displays **local** if the packet was generated by the device's internal Ethernet interface, **em1** or **fxp1**, which connects the Routing Engine with the router's packet-forwarding components.
Name of protocol	Packet's protocol name: **egp**, **gre**, **icmp**, **ipip**, **ospf**, **pim**, **rsvp**, **tcp**, or **udp**.
Packet length	Length of the packet.
Source address	Packet's source address.
Destination address	Packet's destination address and port.

Sample Output

show firewall log

```
user@host>show firewall log
Time       Filter    Action Interface     Protocol  Src Addr      Dest Addr

13:10:12   pfe       D       rlsq0.902     ICMP      192.0.2.2   192.0.2.1

13:10:11   pfe       D       rlsq0.902     ICMP      192.0.2.2   192.0.2.1
```

show firewall log detail

```
user@host> show firewall log detail
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0Name of protocol: TCP, Packet Length: 50824, Source address:
203.0.113.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 1020, Source address: 203.0.113.108:829,
Destination address: 192.168.70.66:513
```

 Copyright © 2018, Juniper Networks, Inc.

```
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 203.0.113.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 203.0.113.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 203.0.113.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 203.0.113.108:829,
Destination address: 192.168.70.66:513
....
```

show firewall log extensive

```
user@host> show firewall log extensive
Time of Log: 2016-01-17 22:16:21 PST, Filter: pfe, Filter action: accept, Name
of interface: xe-0/0/1.0
Name of protocol: UDP, Packet Length: 98, Source address: 203.0.113.1, Destination
 address: 203.0.113.1
 :  00-0F: 00 01 03 ee ee ff 00 01 - 09 22 55 ee 81 00 02 58
 :  10-1F: 08 00 45 00 00 62 00 00 - 00 00 40 11 77 8a 01 00
 :  20-2F: 00 01 02 00 00 01 1c 00 - 1c 00 00 4e 19 83 00 01
 :  30-3F: 02 03 04 05 06 07 08 09 - 0a 0b 0c 0d 0e 0f 10 11
 :  40-4F: 12 13 14 15 16 17 18 19 - 1a 1b 1c 1d 1e 1f 20 21
 :  50-5F: 22 23 24 25 26 27 28 29 - 2a 2b 00 00 00 00 00 00
 :  60-6F: 00 00 00 00 00 00 00 00 - 00 00 00 00 00 00 00 00
 :  70-7F: 00 00 00 00 00 00 00 00 - 00 00 00 00 00 00 00 00
```

show firewall prefix-action-stats

List of Syntax

Syntax (filter-specific mode) on page 2456
Syntax (term-specific mode) on page 2456

Syntax (filter-specific mode)

show firewall prefix-action-stats filter *filter-name* prefix-action *prefix-action-name*
<from *number* to *number*>
<logical-system (*logical-system-name* | all)>

Syntax (term-specific mode)

show firewall prefix-action-stats filter *filter-name* prefix-action *prefix-action-name-term-name*
<from *number* to *number*>
<logical-system (*logical-system-name* | all)>

Release Information

Command introduced before Junos OS Release 7.4.
logical-system option introduced in Junos OS Release 9.3.

Description

Display prefix action statistics about configured firewall filters.

If you clear statistics for firewall filters that are applied to MPCs and that also use the **prefix-action** action on matched packets, wait at least 5 seconds before you enter the **show firewall prefix-action-stats** command. A 5-second pause between issuing the **clear firewall** and **show firewall prefix-action-stats** commands avoids a possible timeout of the **show firewall prefix-action-stats** command.

By default, policers operate in *term-specific* mode.

See *Filter-Specific Policer Overview* for information about how to configure policers in *filter-specific* mode.

Options

filter *filter-name*—Name of a filter.

prefix-action *prefix-action-name*—Name of a prefix action.

from *number* to *number*—(Optional) Starting and ending counter or policer.

logical-system (*logical-system-name* | all)—(Optional) Perform this operation on all logical systems or on a particular system.

Required Privilege Level

view

Related Documentation

- clear firewall on page 2442

List of Sample Output

show firewall prefix-action-stats on page 2457

Output Fields

Table 226 on page 2457 lists the output fields for the **show firewall prefix-action-stats** command. Output fields are listed in the approximate order in which they appear.

 Copyright © 2018, Juniper Networks, Inc.

Table 226: show firewall prefix-action-stats Output Fields

Field Name	Field Description
Filter	Filter name. Filters configured for logical systems include the name of the filter prefixed with the two underscore characters (__) and the name of the logical system (for example, **__ls1/filter1**).

Sample Output

The following sample output assumes that the policer *act1* is in term mode and that there is a term named *term1* configured in the firewall filter *test.*

show firewall prefix-action-stats

```
user@host> show firewall prefix-action-stats filter test prefix-action act1-term1 from 0 to 9
Filter: test
Counters:
Name                        Bytes              Packets
act1-0                          0                    0
act1-1                          0                    0
act1-2                          0                    0
act1-3                          0                    0
act1-4                          0                    0
act1-5                          0                    0
act1-6                          0                    0
act1-7                          0                    0
act1-8                          0                    0
act1-9                          0                    0
Policers:
Name                        Bytes              Packets
act1-0                          0                    0
act1-1                          0                    0
act1-2                          0                    0
act1-3                          0                    0
act1-4                          0                    0
act1-5                          0                    0
act1-6                          0                    0
act1-7                          0                    0
act1-8                          0                    0
act1-9                          0                    0
```

Copyright © 2018, Juniper Networks, Inc.

show firewall templates-in-use

Syntax show firewall templates-in-use

Release Information Command introduced in Junos OS Release 12.3.

Description Display the names of configured filter templates that are currently in use by dynamic subscribers and the number of times each template is referenced.

Required Privilege Level view

Related Documentation

- clear firewall on page 2442
- show firewall log on page 2453

List of Sample Output show firewall templates-in-use on page 2459

Output Fields Table 227 on page 2458 lists the output fields for the **show firewall templates-in-use** command. Output fields are listed in the approximate order in which they appear.

Table 227: show firewall templates-in-use Output Fields

Field Name	Field Description
Filter Template	Name of a filter that has been configured using the **filter** statement at either the **[edit firewall]** or **[edit dynamic-profiles *profile-name* firewall]** hierarchy and is being used as a template for dynamic subscriber filtering.
Reference Count	Number of times the filter has been referenced by subscribers accessing the network.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show firewall templates-in-use

```
user@host> show firewall templates-in-use
                                Dynamic Subscribers Reference Counts
Filter Template
                   Reference Count
---------------
                   ---------------
egressFilter
                         10
ingressFilter
                         10
dfilter
                          5
dfilter-pol
                          5
```

Copyright © 2018, Juniper Networks, Inc.

show policer

Syntax

show policer
<detail>
<*policer-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Option **detail** introduced in Junos OS Release 12.3.

Description

Display the number of policed packets for a given policer or an aggregate policer. An aggregate policer is an aggregate of different policers on the same logical interface.

Options

none—Display the number of policed packets for all configured policers.

detail—(Optional) Display enhanced statistics and counters for policers.

policer-name—(Optional) Display the number of policed packets for the specified policer.

Required Privilege Level

view

List of Sample Output

show policer (MX Series) on page 2461
show policer (non MX Series Router) on page 2461
show policer (Aggregate Policer, non MX Series Router) on page 2461
show policer detail on page 2462

Output Fields

Table 228 on page 2460 lists the output fields for the **show policer** command. Output fields are listed in the approximate order in which they appear.

Table 228: show policer Output Fields

Field Name	Field Description
Name	Name of the policer.
Bytes	• (For two-color policers on MX Series routers, and for hierarchical policers on MS-DPC, MIC, and MPC interfaces on MX Series routers)—Total number of bytes policed by the specified policer. For other combinations of policer type, device, and line card type, this field is blank. • (T Series and M10i)—Not applicable. The Bytes information is not displayed.
Packets	Total number of packets policed by the specified policer.

Copyright © 2018, Juniper Networks, Inc.

Table 228: show policer Output Fields (continued)

Field Name	Field Description
Policer detail	• OOS packet statistics for packets that are marked out-of-specification by the policer. Changes to all packets that have out-of-specification actions, such as discard, color marking, or forwarding-class, are included in this counter. • Offered packet statistics for traffic subjected to policing. • Transmitted packet statistics for traffic that is not discarded by the policer. When the policer action is discard, the statistics are the same as the within-specification statistics; when the policer action is non-discard (loss-priority or forwarding-class), the statistics are included in this counter.

Sample Output

show policer (MX Series)

```
user@host> show policer
Policers:
Name                                              Bytes              Packets
__default_arp_policer__                          314520                 5242
pol-2M-ge-1/2/0.1-inet-i                       10372300               103723
pol-2M-ge-1/2/0.1-inet6-i                       7727800                77278
pol-2M-ge-1/2/0.1-mpls-i                        7070336                67984
pol-2M-ge-1/2/0.1001-vpls-i                    65153700               651537
pol-2M-ge-1/2/0.2001-vpls-i                    65180900               651809
pol-2M-ge-1/2/0.3001-ccc-i                     62202144               647939
```

show policer (non MX Series Router)

```
user@host> show policer
Policers:

Name                                              Bytes              Packets
__default_arp_policer__                              NA                 5242
pol-2M-ge-1/2/0.1-inet-i                             NA               103723
pol-2M-ge-1/2/0.1-inet6-i                            NA                77278
pol-2M-ge-1/2/0.1-mpls-i                             NA                67984
pol-2M-ge-1/2/0.1001-vpls-i                          NA               651537
pol-2M-ge-1/2/0.2001-vpls-i                          NA               651809
pol-2M-ge-1/2/0.3001-ccc-i                           NA               647939
```

show policer (Aggregate Policer, non MX Series Router)

```
user@host> show policer
Policers:
Name                                              Bytes              Packets
__default_arp_policer__                              NA                    0
P1-ae0.0-log_int-o                                   NA                    0
P2-ge-7/0/2.0-inet-o                                 NA                    0
P2-ge-7/0/2.0-inet6-o                                NA                    0
__policer_tmpl__-term                                NA                    0
__policer_tmpl__-fc0                                 NA                    0
__policer_tmpl__-fc0                                 NA                    0
__policer_tmpl__-fc1                                 NA                    0
__policer_tmpl__-fc0                                 NA                    0
```

```
__policer_tmpl__-fc1                                    NA             0
__policer_tmpl__-fc2                                    NA             0
__policer_tmpl__-fc0                                    NA             0
__policer_tmpl__-fc1                                    NA             0
__policer_tmpl__-fc2                                    NA             0
__policer_tmpl__-fc3                                    NA             0
```

show policer detail

```
user@host> show policer detail
Policers:
Name                                                 Bytes        Packets
__default_arp_policer__
    OOS                                                  0              0
    Offered                                              0            496
    Transmitted                                          0            496
P1-xe-1/0/0.0-inet-i
    OOS                                                  0          11329
    Offered                                              0         111188
    Transmitted                                          0          99859
```

 Copyright © 2018, Juniper Networks, Inc.

CHAPTER 29

Layer 2 Bridging and Switching Operational Commands

Copyright © 2018, Juniper Networks, Inc.

clear bridge mac-table

Syntax

clear bridge mac-table
<bridge-domain (all | *bridge-domain-name*)>
<instance *instance-name*>
<interface *interface-name*>
<learning-vlan id (all-vlan | *learning-vlan-id*)>
<*mac-address*>

Release Information

Command introduced in Junos OS Release 8.4.

Description

(MX Series routers only) Clear learned Layer 2 address information from the media access control (MAC) address table.

Options

none—Clear all learned Layer 2 address information from the MAC address table.

bridge-domain (all | *bridge-domain-name*)—(Optional) Clear learned Layer 2 MAC addresses for all bridging domains or for the specified bridging domain.

instance *instance-name*—(Optional) Clear learned Layer 2 MAC addresses for the specified routing instance.

interface *interface-name*—(Optional) Clear learned Layer 2 MAC addresses for the specified interface.

learning-vlan-id (all-vlan | *learning-vlan-id*)—(Optional) Clears learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.

mac-address—(Optional) Clear the specified learned Layer 2 address from the MAC address table.

Required Privilege Level

clear

List of Sample Output

clear bridge mac-table on page 2464

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear bridge mac-table

```
user@host> clear bridge mac-table
```

 Copyright © 2018, Juniper Networks, Inc.

clear error bpdu interface

List of Syntax
MX Series on page 2465
QFX Series, EX Series, NFX Series on page 2465

MX Series
clear error bpdu interface *interface-name*

QFX Series, EX Series, NFX Series
clear error bpdu interface (all | *interface-name*)

Release Information
Command introduced in Junos OS Release 9.4.
Command introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
Command supports **all** option in Junos OS Release 15.1 for EX Series switches.

Description
Clear a bridge protocol data unit (BPDU) error condition caused by the detection of a possible bridging loop from Spanning Tree Protocol (STP) operation.

Required Privilege Level
clear

Related Documentation

- *Configuring BPDU Protection for Spanning-Tree Instance Interfaces*
- *Configuring BPDU Protection on All Edge Ports*
- *Unblocking a Switch Interface That Receives BPDUs in Error (CLI Procedure)*

List of Sample Output
clear error bpdu interface on page 2465

Output Fields
When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear error bpdu interface

```
user@host> clear error bpdu interface ge-1/1/1
```

clear error mac-rewrite

Syntax

clear error mac-rewrite
<interface *interface-name*>

Release Information

Command introduced in Junos OS Release 9.1.

Description

Clear a MAC rewrite error condition caused by the reception of tunneled protocol packets on an interface with Layer 2 protocol tunneling enabled.

On interfaces with L2PT configured, customer-facing ports should not receive packets with the L2PT MAC address as the destination address unless there is a network tolopogy or configuration error. Any such interface receiving an L2PT packet becomes “Disabled”, and must subsequently be re-enabled by clearing the error with this command.

Options

interface *interface-name*—(Optional) Clear the MAC rewrite error condition for the specified interface.

Required Privilege Level

clear

Related Documentation

- *Understanding Layer 2 Protocol Tunneling Through a Network*
- *Configuring Layer 2 Protocol Tunneling on EX Series Switches with ELS Support (CLI Procedure)*
- *Clearing a MAC Rewrite Error on an Interface with Layer 2 Protocol Tunneling*
- show mac-rewrite interface on page 2492

List of Sample Output

clear error mac-rewrite interface on page 2466

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear error mac-rewrite interface

```
user@host> clear error mac-rewrite interface ge-1/0/1
```


Copyright © 2018, Juniper Networks, Inc.

show bridge domain

Syntax

show bridge domain
<brief | detail | extensive>
<bridge-domain (all | *domain-name*)>
<instance *instance-name*>
<operational>

Release Information

Command introduced in Junos OS Release 8.4.

Description

(MX Series routers only) Display bridge domain information.

Options

none—Display information for all bridge domains.

brief | detail | extensive—(Optional) Display the specified level of output.

bridge-domain (all | *domain-name*)— (Optional) Display information about all bridge domains or the specified bridge domain.

instance *instance-name*—(Optional) Display information for the specified routing instance.

operational—(Optional) Display information for the operational routing instances.

Required Privilege Level

view

List of Sample Output

show bridge domain on page 2467
show bridge domain brief on page 2467
show bridge domain detail on page 2468

Sample Output

show bridge domain

```
user@host> show bridge domain
Instance                    Bridging Domain         Type
            Primary Table                                           Active
vs1                         vlan100                 bridge
            bridge.0                                                2
vs1                         vlan200                 bridge
            bridge.0                                                0
```

show bridge domain brief

```
user@host> show bridge domain brief
Instance                    Bridging Domain         Type
            Primary Table                                           Active
vs1                         vlan100                 bridge
            bridge.0                                                2
vs1                         vlan200                 bridge
            bridge.0                                                0
```

Copyright © 2018, Juniper Networks, Inc.

show bridge domain detail

```
user@host> show bridge domain detail
Routing Instance:vs1
  Bridging Domain:vlan100
  Router ID: 0.0.0.0
  Type: bridge              State: Active
  Interfaces:
    ge-11/0/3.0
    ge-11/1/4.100
    ge-11/1/1.100
    ge-11/1/0.100
    xe-10/2/0.100
    xe-10/0/0.100
  Tables:
     bridge.0               : 2 macs (2 active)
Routing Instance:vs1
  Bridging Domain:vlan200
  Router ID: 0.0.0.0
  Type: bridge              State: Active
  Interfaces:
    ge-11/1/0.200
    ge-11/1/1.200
    ge-11/1/4.200
    xe-10/0/0.200
    xe-10/2/0.200
  Tables:
     bridge.0               : 0 macs (0 active)
```

 Copyright © 2018, Juniper Networks, Inc.

show bridge flood

Syntax

show bridge flood
<brief | detail | extensive>
<bridge-domain *domain-name*>
<event-queue>
<instance *instance-name*>
<route (all-ce-flood | all ve-flood | alt-root-flood | bd-flood | mlp-flood | re-flood)>

Release Information

Command introduced in Junos OS Release 8.4.

Description

(MX Series routers only) Display bridging flooding information.

Options

none—Display all bridging flooding information for all bridging domains.

brief | detail | extensive—(Optional) Display the specified level of output.

bridge-domain *domain-name*—(Optional) Display bridging flooding information for the specified bridge domain.

event-queue—(Optional) Display the queue of pending bridge flood events.

instance *instance-name*—(Optional) Display bridging flooding information for the specified routing instance.

route (all-ce-flood | all ve-flood | alt-root-flood | bd-flood | mlp-flood | re-flood)—(Optional) Display the following:

- **all-ce-flood**—Display the route for flooding traffic to all customer edge routers if **no-local-switching** is enabled.
- **all-ve-flood**—Display the route for flooding traffic to all VPLS edge routers if **no-local-switching** is enabled.
- **alt-root-flood**—Display the Spanning Tree Protocol (STP) alt-root flooding route used for the interface.
- **bd-flood**—Display the route for flooding traffic of a bridge domain if **no-local-switching** is not enabled.
- **mlp-flood**—Display the route for flooding traffic to MAC learning chips.
- **re-flood**—Display the route for Routing Engine flooding to all interfaces.

Required Privilege Level

view

List of Sample Output

show bridge flood on page 2470
show bridge flood brief on page 2470
show bridge flood detail on page 2470
show bridge flood extensive on page 2471

Output Fields to be provided

Sample Output

show bridge flood

```
user@host> show bridge flood
Name: __juniper_private1__
CEs: 0
VEs: 0
Flood Routes:
  Prefix     Type          Owner                 NhType          NhIndex
  0x36/16    MLP_FLOOD     __vs1+vlan100__       flood           426
  0x3a/16    MLP_FLOOD     __vs1+vlan200__       flood           428
Name: vs1::vlan100
CEs: 6
VEs: 0
Flood Routes:
  Prefix     Type          Owner                 NhType          NhIndex
  0x35/16    ALL_FLOOD     __vs1+vlan100__       flood           425
  0x35/16    RE_FLOOD      __vs1+vlan100__       flood           425
  0x3780/17 ALT_ROOT_RT    ge-11/0/3.0           flood           425
  0x3b80/17 ALT_ROOT_RT    ge-11/1/4.100         flood           425
  0x3c80/17 ALT_ROOT_RT    ge-11/1/1.100         flood           425
  0x3d80/17 ALT_ROOT_RT    ge-11/1/0.100         flood           425
  0x3e80/17 ALT_ROOT_RT    xe-10/2/0.100         flood           425
  0x3f80/17 ALT_ROOT_RT    xe-10/0/0.100         flood           425
Name: vs1::vlan200
CEs: 5
VEs: 0
Flood Routes:
  Prefix     Type          Owner                 NhType          NhIndex
  0x39/16    ALL_FLOOD     __vs1+vlan200__       flood           427
  0x39/16    RE_FLOOD      __vs1+vlan200__       flood           427
  0x4180/17 ALT_ROOT_RT    ge-11/1/0.200         flood           427
  0x4080/17 ALT_ROOT_RT    ge-11/1/1.200         flood           427
  0x4280/17 ALT_ROOT_RT    ge-11/1/4.200         flood           427
  0x4480/17 ALT_ROOT_RT    xe-10/0/0.200         flood           427
  0x4380/17 ALT_ROOT_RT    xe-10/2/0.200         flood           427
```

show bridge flood brief

```
user@host> show bridge flood brief
Name                 Active CEs         Active VEs
__juniper_private1__ 0                  0
vs1::vlan100         6                  0
vs1::vlan200         5                  0
```

show bridge flood detail

```
user@host> show bridge flood detail
Name: __juniper_private1__
CEs: 0
VEs: 0
Flood Routes:
  Prefix     Type          Owner                 NhType          NhIndex
  0x36/16    MLP_FLOOD     __vs1+vlan100__       flood           426
  0x3a/16    MLP_FLOOD     __vs1+vlan200__       flood           428
Name: vs1::vlan100
```

 Copyright © 2018, Juniper Networks, Inc.

```
CEs: 6
VEs: 0
Flood Routes:
  Prefix     Type           Owner                 NhType          NhIndex
  0x35/16    ALL_FLOOD      __vs1+vlan100__       flood           425
  0x35/16    RE_FLOOD       __vs1+vlan100__       flood           425
  0x3780/17 ALT_ROOT_RT    ge-11/0/3.0           flood           425
  0x3b80/17 ALT_ROOT_RT    ge-11/1/4.100         flood           425
  0x3c80/17 ALT_ROOT_RT    ge-11/1/1.100         flood           425
  0x3d80/17 ALT_ROOT_RT    ge-11/1/0.100         flood           425
  0x3e80/17 ALT_ROOT_RT    xe-10/2/0.100         flood           425
  0x3f80/17 ALT_ROOT_RT    xe-10/0/0.100         flood           425
Name: vs1::vlan200
CEs: 5
VEs: 0
Flood Routes:
  Prefix     Type           Owner                 NhType          NhIndex
  0x39/16    ALL_FLOOD      __vs1+vlan200__       flood           427
  0x39/16    RE_FLOOD       __vs1+vlan200__       flood           427
  0x4180/17 ALT_ROOT_RT    ge-11/1/0.200         flood           427
  0x4080/17 ALT_ROOT_RT    ge-11/1/1.200         flood           427
  0x4280/17 ALT_ROOT_RT    ge-11/1/4.200         flood           427
  0x4480/17 ALT_ROOT_RT    xe-10/0/0.200         flood           427
  0x4380/17 ALT_ROOT_RT    xe-10/2/0.200         flood           427
```

show bridge flood extensive

```
user@host> show bridge flood extensive
Name: __juniper_private1__
CEs: 0
VEs: 0
  Flood route prefix: 0x36/16
  Flood route type: MLP_FLOOD
  Flood route owner: __vs1+vlan100__
  Nexthop type: flood
  Nexthop index: 426
    Interfaces Flooding to:
    Name              Type          NhType          Index
    lc-11/0/0.32769  LC
    lc-10/2/0.32769  LC
    lc-10/0/0.32769  LC
    lc-11/1/0.32769  LC

  Flood route prefix: 0x3a/16
  Flood route type: MLP_FLOOD
  Flood route owner: __vs1+vlan200__
  Nexthop type: flood
  Nexthop index: 428
    Interfaces Flooding to:
    Name              Type          NhType          Index
    lc-10/0/0.32769  LC
    lc-10/2/0.32769  LC
    lc-11/1/0.32769  LC
Name: vs1::vlan100
CEs: 6
VEs: 0

  Flood route prefix: 0x35/16
  Flood route type: ALL_FLOOD
  Flood route owner: __vs1+vlan100__
  Nexthop type: flood
```

```
Nexthop index: 425
  Interfaces Flooding to:
  Name             Type          NhType          Index
  ge-11/0/3.0      CE
  ge-11/1/4.100    CE
  ge-11/1/1.100    CE
  ge-11/1/0.100    CE
  xe-10/2/0.100    CE
  xe-10/0/0.100    CE

Flood route prefix: 0x35/16
Flood route type: RE_FLOOD
Flood route owner: __vs1+vlan100__
Nexthop type: flood
Nexthop index: 425
  Interfaces Flooding to:
  Name             Type          NhType          Index
  ge-11/0/3.0      CE
  ge-11/1/4.100    CE
  ge-11/1/1.100    CE
  ge-11/1/0.100    CE
  xe-10/2/0.100    CE
  xe-10/0/0.100    CE

Flood route prefix: 0x3780/17
Flood route type: ALT_ROOT_RT
Flood route owner: ge-11/0/3.0
Nexthop type: flood
Nexthop index: 425
  Interfaces Flooding to:
  Name             Type          NhType          Index
  ge-11/0/3.0      CE
  ge-11/1/4.100    CE
  ge-11/1/1.100    CE
  ge-11/1/0.100    CE
  xe-10/2/0.100    CE
  xe-10/0/0.100    CE

Flood route prefix: 0x3b80/17
Flood route type: ALT_ROOT_RT
Flood route owner: ge-11/1/4.100
Nexthop type: flood
Nexthop index: 425
  Interfaces Flooding to:
  Name             Type          NhType          Index
  ge-11/0/3.0      CE
  ge-11/1/4.100    CE
  ge-11/1/1.100    CE
  ge-11/1/0.100    CE
  xe-10/2/0.100    CE
  xe-10/0/0.100    CE

Flood route prefix: 0x3c80/17
Flood route type: ALT_ROOT_RT
Flood route owner: ge-11/1/1.100
Nexthop type: flood
Nexthop index: 425
  Interfaces Flooding to:
  Name             Type          NhType          Index
  ge-11/0/3.0      CE
  ge-11/1/4.100    CE
```

 Copyright © 2018, Juniper Networks, Inc.

```
      ge-11/1/1.100    CE
      ge-11/1/0.100    CE
      xe-10/2/0.100    CE
      xe-10/0/0.100    CE

    Flood route prefix: 0x3d80/17
    Flood route type: ALT_ROOT_RT
    Flood route owner: ge-11/1/0.100
    Nexthop type: flood
    Nexthop index: 425
      Interfaces Flooding to:
      Name             Type          NhType          Index
      ge-11/0/3.0      CE
      ge-11/1/4.100    CE
      ge-11/1/1.100    CE
      ge-11/1/0.100    CE
      xe-10/2/0.100    CE
      xe-10/0/0.100    CE

    Flood route prefix: 0x3e80/17
    Flood route type: ALT_ROOT_RT
    Flood route owner: xe-10/2/0.100
    Nexthop type: flood
    Nexthop index: 425
      Interfaces Flooding to:
      Name             Type          NhType          Index
      ge-11/0/3.0      CE
      ge-11/1/4.100    CE
      ge-11/1/1.100    CE
      ge-11/1/0.100    CE
      xe-10/2/0.100    CE
      xe-10/0/0.100    CE

    Flood route prefix: 0x3f80/17
    Flood route type: ALT_ROOT_RT
    Flood route owner: xe-10/0/0.100
    Nexthop type: flood
    Nexthop index: 425
      Interfaces Flooding to:
      Name             Type          NhType          Index
      ge-11/0/3.0      CE
      ge-11/1/4.100    CE
      ge-11/1/1.100    CE
      ge-11/1/0.100    CE
      xe-10/2/0.100    CE
      xe-10/0/0.100    CE
  Name: vs1::vlan200
  CEs: 5
  VEs: 0

    Flood route prefix: 0x39/16
    Flood route type: ALL_FLOOD
    Flood route owner: __vs1+vlan200__
    Nexthop type: flood
    Nexthop index: 427
      Interfaces Flooding to:
      Name             Type          NhType          Index
      ge-11/1/0.200    CE
      ge-11/1/1.200    CE
      ge-11/1/4.200    CE
      xe-10/0/0.200    CE
```

```
  xe-10/2/0.200    CE

Flood route prefix: 0x39/16
Flood route type: RE_FLOOD
Flood route owner: __vs1+vlan200__
Nexthop type: flood
Nexthop index: 427
  Interfaces Flooding to:
  Name             Type          NhType          Index
  ge-11/1/0.200    CE
  ge-11/1/1.200    CE
  ge-11/1/4.200    CE
  xe-10/0/0.200    CE
  xe-10/2/0.200    CE

Flood route prefix: 0x4180/17
Flood route type: ALT_ROOT_RT
Flood route owner: ge-11/1/0.200
Nexthop type: flood
Nexthop index: 427
  Interfaces Flooding to:
  Name             Type          NhType          Index
  ge-11/1/0.200    CE
  ge-11/1/1.200    CE
  ge-11/1/4.200    CE
  xe-10/0/0.200    CE
  xe-10/2/0.200    CE

Flood route prefix: 0x4080/17
Flood route type: ALT_ROOT_RT
Flood route owner: ge-11/1/1.200
Nexthop type: flood
Nexthop index: 427
  Interfaces Flooding to:
  Name             Type          NhType          Index
  ge-11/1/0.200    CE
  ge-11/1/1.200    CE
  ge-11/1/4.200    CE
  xe-10/0/0.200    CE
  xe-10/2/0.200    CE

Flood route prefix: 0x4280/17
Flood route type: ALT_ROOT_RT
Flood route owner: ge-11/1/4.200
Nexthop type: flood
Nexthop index: 427
  Interfaces Flooding to:
  Name             Type          NhType          Index
  ge-11/1/0.200    CE
  ge-11/1/1.200    CE
  ge-11/1/4.200    CE
  xe-10/0/0.200    CE
  xe-10/2/0.200    CE

Flood route prefix: 0x4480/17
Flood route type: ALT_ROOT_RT
Flood route owner: xe-10/0/0.200
Nexthop type: flood
Nexthop index: 427
  Interfaces Flooding to:
  Name             Type          NhType          Index
```

 Copyright © 2018, Juniper Networks, Inc.

```
  ge-11/1/0.200    CE
  ge-11/1/1.200    CE
  ge-11/1/4.200    CE
  xe-10/0/0.200    CE
  xe-10/2/0.200    CE

Flood route prefix: 0x4380/17
Flood route type: ALT_ROOT_RT
Flood route owner: xe-10/2/0.200
Nexthop type: flood
Nexthop index: 427
  Interfaces Flooding to:
  Name             Type          NhType          Index
  ge-11/1/0.200    CE
  ge-11/1/1.200    CE
  ge-11/1/4.200    CE
  xe-10/0/0.200    CE
  xe-10/2/0.200    CE
```

show bridge mac-table

Syntax

show bridge mac-table
<age>
<brief | count | detail | extensive>
<bridge-domain (all | ***bridge-domain-name*****)>**
<global-count>
<instance ***instance-name*****>**
<interface ***interface-name*****>**
<***mac-address*****>**
<instance ***instance-name*****>**
<vlan-id (all-vlan | ***vlan-id*****)>**

Release Information

Command introduced in Junos OS Release 8.4.
Command introduced in Junos OS Release 15.1
Support for PBB-EVPN instance added in Junos OS Release 16.1
MAC Flag P to indicate a MAC Pinned interface introduced in Junos OS 16.2

Description

(MX Series routers only) Display Layer 2 MAC address information.

Options

none—Display all learned Layer 2 MAC address information.

age— (Optional) Display age of a single mac-address.

brief | count | detail | extensive—(Optional) Display the specified level of output.

bridge-domain (all | ***bridge-domain-name*****)**—(Optional) Display learned Layer 2 MAC addresses for all bridging domains or for the specified bridging domain.

global-count—(Optional) Display the total number of learned Layer 2 MAC addresses on the system.

instance ***instance-name***—(Optional) Display learned Layer 2 MAC addresses for the specified routing instance.

interface ***interface-name***—(Optional) Display learned Layer 2 MAC addresses for the specified interface.

mac-address—(Optional) Display the specified learned Layer 2 MAC address information.

vlan-id (all-vlan | ***vlan-id*****)**—(Optional) Display learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.

Additional Information

When Layer 2 protocol tunneling is enabled, the tunneling MAC address 01:00:0c:cd:cd:d0 is installed in the MAC table. When the Cisco Discovery Protocol (CDP), Spanning Tree Protocol (STP), or VLAN Trunk Protocol (VTP) is configured for Layer 2 protocol tunneling on an interface, the corresponding protocol MAC address is installed in the MAC table.

 Copyright © 2018, Juniper Networks, Inc.

Required Privilege Level view

List of Sample Output

show bridge mac-table on page 2478
show bridge mac-table (with Layer 2 Services over GRE Interfaces) on page 2478
show bridge mac-table (with VXLAN enabled) on page 2479
show bridge mac-table age (for GE interface) on page 2479
show bridge mac-table age (for AE interface) on page 2479
show bridge mac-table count on page 2479
show bridge mac-table detail on page 2480
show bridge mac-table instance pbb-evpn on page 2480
show bridge mac-table on page 2480

Output Fields Table 229 on page 2477 describes the output fields for the **show bridge mac-table** command. Output fields are listed in the approximate order in which they appear.

Table 229: show bridge mac-table Output Fields

Field Name	Field Description
Age	Age of a single mac-address.
Routing instance	Name of the routing instance.
Bridging domain	Name of the bridging domain.
MAC address	MAC address or addresses learned on a logical interface.
MAC flags	Status of MAC address learning properties for each interface: • **S**—Static MAC address is configured. • **D**—Dynamic MAC address is configured. • **L**—Locally learned MAC address is configured. • **C**—Control MAC address is configured. • **SE**—MAC accounting is enabled. • **NM**—Non-configured MAC. • **R**—Remote PE MAC address is configured. • **P**—MAC Pinned interface is configured
Logical interface	Name of the logical interface.
MAC count	Number of MAC addresses learned on the specific routing instance or interface.
Learning interface	Name of the logical interface on which the MAC address was learned.
Learning VLAN	VLAN ID of the routing instance or bridge domain in which the MAC address was learned.
VXLAN ID/VXLAN	VXLAN Network Identifier (VNI).

Copyright © 2018, Juniper Networks, Inc.

Table 229: show bridge mac-table Output Fields (continued)

Field Name	Field Description
Layer 2 flags	Debugging flags signifying that the MAC address is present in various lists.
Epoch	Spanning Tree Protocol epoch number identifying when the MAC address was learned. Used for debugging.
Sequence number	Sequence number assigned to this MAC address. Used for debugging.
Learning mask	Mask of the Packet Forwarding Engines where this MAC address was learned. Used for debugging.
IPC generation	Creation time of the logical interface when this MAC address was learned. Used for debugging.

Sample Output

show bridge mac-table

```
user@host> show bridge mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
           SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : default-switch
 Bridging domain : test1, VLAN : 1
   MAC                 MAC      Logical          NH     RTR
   address             flags    interface        Index  ID
   01:00:0c:cc:cc:cc   S,NM     NULL
   01:00:0c:cc:cc:cd   S,NM     NULL
   01:00:0c:cd:cd:d0   S,NM     NULL
   64:87:88:6a:17:d0   D        ae0.1
   64:87:88:6a:17:f0   D        ae0.1
```

show bridge mac-table (with Layer 2 Services over GRE Interfaces)

```
user@host> show bridge mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned
           SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : default-switch
 Bridging domain : vlan-1, VLAN : 1
   MAC                 MAC      Logical
   address             flags    interface
   00:01:01:00:01:f7   D,SE     gr-1/2/10.0
   00:03:00:32:01:f7   D,SE     gr-1/2/10.0
   00:00:21:11:11:10   DL       ge-1/0/0.0
   00:00:21:11:11:11   DL       ge-1/1/0.0


Routing instance : default-switch
 Bridging domain : vlan-2, VLAN : 2
   MAC                 MAC      Logical
   address             flags    interface
   00:02:01:33:01:f7   D,SE     gr-1/2/10.1
```

 Copyright © 2018, Juniper Networks, Inc.

```
   00:00:21:11:21:10   DL       ge-1/0/0.1
   00:00:21:11:21:11   DL       ge-1/1/0.1
```

show bridge mac-table (with VXLAN enabled)

```
user@host> show bridge mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned
           SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : default-switch
 Bridging domain : vlan-1, VLAN : 1
  VXLAN: Id : 100, Multicast group: 233.252.0.1
   MAC                 MAC      Logical
   address             flags    interface
   00:01:01:00:01:f7   D,SE     vtep.1052010
   00:03:00:32:01:f7   D,SE     vtep.1052011
   00:00:21:11:11:10   DL       ge-1/0/0.0
   00:00:21:11:11:11   DL       ge-1/1/0.0


Routing instance : default-switch
 Bridging domain : vlan-2, VLAN : 2, VXLAN : 200
  VXLAN: Id : 200, Multicast group: 233.252.0.2
   MAC                 MAC      Logical
   address             flags    interface
   00:02:01:33:01:f7   D,SE     vtep.1052010
   00:04:00:14:01:f7   D,SE     vtep.1052011
   00:00:21:11:21:10   DL       ge-1/0/0.1
   00:00:21:11:21:11   DL       ge-1/1/0.1
```

show bridge mac-table age (for GE interface)

```
user@host> show vpls mac-table age 00:02:03:aa:bb:1a instance vpls_instance_1
MAC Entry Age information
Current Age: 4 seconds
```

show bridge mac-table age (for AE interface)

```
user@host> show vpls mac-table age 00:02:03:aa:bb:1a instance vpls_instance_1
MAC Entry Age information
Current Age on FPC1: 102 seconds
Current Age on FPC2: 94 seconds
```

show bridge mac-table count

```
user@host> show bridge mac-table count
2 MAC address learned in routing instance vs1 bridge domain vlan100

  MAC address count per interface within routing instance:
    Logical interface        MAC count
    ge-11/0/3.0                      1
    ge-11/1/4.100                    0
    ge-11/1/1.100                    0
    ge-11/1/0.100                    0
    xe-10/2/0.100                    1
    xe-10/0/0.100                    0

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID            MAC count
```

Copyright © 2018, Juniper Networks, Inc.

```
                0                    2

0 MAC address learned in routing instance vs1 bridge domain vlan200

  MAC address count per interface within routing instance:
    Logical interface        MAC count
    ge-11/1/0.200                    0
    ge-11/1/1.200                    0
    ge-11/1/4.200                    0
    xe-10/0/0.200                    0
    xe-10/2/0.200                    0

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID            MAC count
                0                    0
```

show bridge mac-table detail

```
user@host> show bridge mac-table detail
MAC address: 00:00:00:19:1c:db
  Routing instance: vs1
   Bridging domain: vlan100
   Learning interface: ge-11/0/3.0     Learning VLAN: 0
   Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
   Epoch: 4                            Sequence number: 0
   Learning mask: 0x800                IPC generation: 0

MAC address: 00:00:00:59:3a:2f
  Routing instance: vs1
   Bridging domain: vlan100
   Learning interface: xe-10/2/0.100   Learning VLAN: 0
   Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
   Epoch: 7                            Sequence number: 0
   Learning mask: 0x400                IPC generation: 0
```

show bridge mac-table instance pbb-evpn

```
user@host> show bridge mac-table instance pbb-evpn
Routing instance : pbb-evpn
 Bridging domain : isid-bd10000, ISID : 10000
   MAC                       MAC      Logical          NH     RTR
   address                  flags      interface        Index  ID
   00:19:e2:b0:76:eb    D         cbp.1000
   aa:bb:cc:dd:ee:f2    DC                         1048576 1048576
   aa:bb:cc:dd:ee:f3    DC                         1048575 1048575
```

show bridge mac-table

```
user@host>run show bridge mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
O -OVSDB MAC, SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC,
P -Pinned MAC)

Routing instance : VS-541
Bridging domain : 541, VLAN : 541
MAC MAC Logical NH RTR
address flags interface Index ID
00:00:01:00:00:01 DPRC xe-0/0/3.0
00:00:02:00:00:01 DP xe-0/0/3.0
```

 Copyright © 2018, Juniper Networks, Inc.

show bridge statistics

Syntax

show bridge statistics
<bridge-domain *domain-name*>
<instance *instance-name*>

Release Information

Command introduced in Junos OS Release 8.4.

Description

(MX Series routers only) Display bridge statistics.

Options

none—Display bridge statistics for all bridge domains in all routing instances.

bridge-domain *domain-name*—(Optional) Display statistics for the specified bridge domain.

instance *instance-name*—(Optional) Display statistics for the specified routing instance.

Required Privilege Level

view

List of Sample Output

show bridge statistics on page 2482

Sample Output

show bridge statistics

```
user@host> show bridge statistics
Information for routing instance:

Routing instance : __juniper_private1__
   Index: 1                         Sequence number: 0
   MAC limit: 5000                  MACs learned: 0
   Static MACs learned: 0           Non config Static MACs learned: 0
   Handle: 0x829e800

Information for routing instance:

Routing instance : vs1
 Bridging domain : vlan100
   Index: 3                         Sequence number: 0
   MAC limit: 5120                  MACs learned: 2
   Static MACs learned: 0           Non config Static MACs learned: 0
   Handle: 0x829e400
   Flags: Bridge instance, Config defined, VLAN : 100
   Local interface: ge-11/0/3.0, Index: 79
     Broadcast packets:                      1
     Broadcast bytes  :                     65
     Multicast packets:                      0
     Multicast bytes  :                      0
     Flooded packets  :                      0
     Flooded bytes    :                      0
     Unicast packets  :              358624489
     Unicast bytes    :            23310592305
     Current MAC count:                      1 (Limit 1024)
```

 Copyright © 2018, Juniper Networks, Inc.

```
   Local interface: ge-11/1/4.100, Index: 84
     Broadcast packets:                     0
     Broadcast bytes  :                     0
     Multicast packets:                     0
     Multicast bytes  :                     0
     Flooded packets  :                     0
     Flooded bytes    :                     0
     Unicast packets  :                     0
     Unicast bytes    :                     0
     Current MAC count:                     0 (Limit 1024)
   Local interface: ge-11/1/1.100, Index: 86
     Broadcast packets:                     0
     Broadcast bytes  :                     0
     Multicast packets:                     0
     Multicast bytes  :                     0
     Flooded packets  :                     0
     Flooded bytes    :                     0
     Unicast packets  :                     0
     Unicast bytes    :                     0
     Current MAC count:                     0 (Limit 1024)
   Local interface: ge-11/1/0.100, Index: 87
     Broadcast packets:                     0
     Broadcast bytes  :                     0
     Multicast packets:                     0
     Multicast bytes  :                     0
     Flooded packets  :                     0
     Flooded bytes    :                     0
     Unicast packets  :                     0
     Unicast bytes    :                     0
     Current MAC count:                     0 (Limit 1024)
   Local interface: xe-10/2/0.100, Index: 88
     Broadcast packets:                     0
     Broadcast bytes  :                     0
     Multicast packets:                     0
     Multicast bytes  :                     0
     Flooded packets  :                     0
     Flooded bytes    :                     0
     Unicast packets  :             358627393
     Unicast bytes    :           23310781065
     Current MAC count:                     1 (Limit 1024)
   Local interface: xe-10/0/0.100, Index: 89
     Broadcast packets:                     0
     Broadcast bytes  :                     0
     Multicast packets:                     0
     Multicast bytes  :                     0
     Flooded packets  :                     0
     Flooded bytes    :                     0
     Unicast packets  :                     0
     Unicast bytes    :                     0
     Current MAC count:                     0 (Limit 1024)

Information for routing instance:

Routing instance : vs1
 Bridging domain : vlan200
   Index: 4                        Sequence number: 0
   MAC limit: 5120                 MACs learned: 0
   Static MACs learned: 0          Non config Static MACs learned: 0
   Handle: 0x829e600
   Flags: Bridge instance, Config defined, VLAN : 200
   Local interface: ge-11/1/0.200, Index: 90
```

```
  Broadcast packets:                 0
  Broadcast bytes  :                 0
  Multicast packets:                 0
  Multicast bytes  :                 0
  Flooded packets  :                 0
  Flooded bytes    :                 0
  Unicast packets  :                 0
  Unicast bytes    :                 0
  Current MAC count:                 0 (Limit 1024)
Local interface: ge-11/1/1.200, Index: 91
  Broadcast packets:                 0
  Broadcast bytes  :                 0
  Multicast packets:                 0
  Multicast bytes  :                 0
  Flooded packets  :                 0
  Flooded bytes    :                 0
  Unicast packets  :                 0
  Unicast bytes    :                 0
  Current MAC count:                 0 (Limit 1024)
Local interface: ge-11/1/4.200, Index: 92
  Broadcast packets:                 0
  Broadcast bytes  :                 0
  Multicast packets:                 0
  Multicast bytes  :                 0
  Flooded packets  :                 0
  Flooded bytes    :                 0
  Unicast packets  :                 0
  Unicast bytes    :                 0
  Current MAC count:                 0 (Limit 1024)
Local interface: xe-10/0/0.200, Index: 93
  Broadcast packets:                 0
  Broadcast bytes  :                 0
  Multicast packets:                 0
  Multicast bytes  :                 0
  Flooded packets  :                 0
  Flooded bytes    :                 0
  Unicast packets  :                 0
  Unicast bytes    :                 0
  Current MAC count:                 0 (Limit 1024)
Local interface: xe-10/2/0.200, Index: 94
  Broadcast packets:                 4
  Broadcast bytes  :               260
  Multicast packets:                 0
  Multicast bytes  :                 0
  Flooded packets  :                 0
  Flooded bytes    :                 0
  Unicast packets  :                 0
  Unicast bytes    :                 0
  Current MAC count:                 0 (Limit 1024)
```

 Copyright © 2018, Juniper Networks, Inc.

show l2-learning global-information

Syntax **show l2-learning global-information**

Release Information Command introduced in Junos OS Release 8.4.

Description (MX Series routers only) Display Layer 2 learning process-related information for the entire router.

Options This command has no options.

Required Privilege Level view

List of Sample Output show l2-learning global-information on page 2485

Output Fields Table 230 on page 2485 describes the output fields for the **show l2-learning global-information** command. Output fields are listed in the approximate order in which they appear.

Table 230: show l2-learning global-information Output Fields

Field Name	Field Description
MAC aging interval	Configured timeout interval, in seconds, for all MAC table entries.
MAC learning	Status of MAC learning: **Enabled** or **Disabled**.
MAC statistics	Status of MAC accounting: **Enabled** or **Disabled**.
MAC limit Count	Configured maximum limit on the number of MAC addresses that can be learned.
MAC limit hit flag	Status of the learned MAC limit hit flag: **Enabled** (the learned MAC exceeds the global MAC limit) or **Disabled** (the learned MAC does not exceed the global MAC limit).
MAC packet action drop	Status of action to drop packets after the configured MAC address limit is reached: **Enabled** (packets are dropped) or **Disabled** (packets are forwarded).

Sample Output

show l2-learning global-information

```
user@host> show l2-learning global-information
Global Configuration:

MAC aging interval    : 300
MAC learning          : Enabled
```

```
MAC statistics        : Disabled
MAC limit Count       : 393215
MAC limit hit flag    : Disabled
MAC packet action drop: Disabled
```

 Copyright © 2018, Juniper Networks, Inc.

show l2-learning global-mac-count

Syntax **show l2-learning global-mac-count**

Release Information Command introduced in Junos OS Release 9.3.

Description (MX Series routers only) Display the total number of dynamic and static MAC addresses learned for the entire router.

Options This command has no options.

Required Privilege Level view

List of Sample Output show l2-learning global-mac-count on page 2487

Output Fields Displays the total number of dynamic and static MAC addresses learned for the entire router.

Sample Output

show l2-learning global-mac-count

```
user@host> show l2-learning global-mac-count
100 dynamic and static MAC addresses learned globally
```

Copyright © 2018, Juniper Networks, Inc.

show l2-learning instance

Syntax **show l2-learning instance**

Release Information (MX Series routers only) Command introduced in Junos OS Release 8.4.

Description Display Layer 2 learning properties for all the configured routing instances.

Options This command has no options.

Required Privilege Level view

List of Sample Output show l2-learning instance on page 2489

Output Fields Table 231 on page 2488 describes the output fields for the **show l2-learning instance** command. Output fields are listed in the approximate order in which they appear.

Table 231: show l2-learning instance Output Fields

Field Name	Field Description
Routing Instance	Name of routing instance.
Bridging Domain	Name of bridging domain. On MX Series routers you can use the **show l2-learning instance <extensive>** command option to display the Bridge Service-id information which includes the Config Service ID and the Active Service ID.
Index	Number associated with the routing instance or bridging domain.
Logical System	Name of logical system or **Default** if no logical system is configured.
Routing instance flags	Status of Layer 2 learning properties for each routing instance: • **DL**—MAC learning is disabled. • **SE**—MAC accounting is enabled. • **AD**—Packets are dropped after MAC address limit is reached. • **LH**—The maximum number of MAC addresses has been learned on the routing instance. The routing instance is not able to learn any additional MAC addresses.
MAC limit	Maximum number of MAC addresses that can be learned from each interface in the routing instance or bridging domain.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show l2-learning instance

```
user@host> show l2-learning instance
Information for routing instance:

Routing Instance flags (DL -disable learning, SE -stats enabled,
                        AD -packet action drop, LH -mac limit hit)

Routing              Bridging      Index   Logical       Routing     MAC
Instance             Domain                System        flags       limit
__juniper_private1__                   1   Default                    5000
vs1                  vlan100           3   Default                    5120
vs1                  vlan200           4   Default                    5120
```

show l2-learning interface

Syntax **show l2-learning interface**

Release Information Command introduced in Junos OS Release 8.4.
Added sample output to indicate an EVPN MAC Pinned interface, introduced in Junos OS 16.2R1.

Description (MX Series routers only) Display Layer 2 learning information for all the interfaces.

Options This command has no options.

Required Privilege Level view

List of Sample Output show l2-learning interface on page 2490
show l2 learning-interface on page 2491

Output Fields Table 232 on page 2490 describes the output fields for the **show l2-learning interface** command. Output fields are listed in the approximate order in which they appear.

Table 232: show l2-learning interfaceOutput Fields

Field Name	Field Description
Logical interface	Name of the logical interface.
Index	Index of the interface.
Routing Instance	Number of the routing instance to which the interface belongs.
Interface device	Value of the order in which the Junos OS finds and initializes the interface.
Logical interface flags	Status of Layer 2 learning properties for each interface: • **DL**—MAC learning is disabled. • **SE**—MAC accounting is enabled. • **AD**—Packets are dropped after the MAC interface limit is reached. • **MAC limit**—Maximum number of MAC addresses that can be learned from the interface. • **MP**—MAC Pinning enabled.

Sample Output

show l2-learning interface

```
user@host> showl2-learning interface
Information for interface family:
```


Copyright © 2018, Juniper Networks, Inc.

```
Logical Interface flags (DL -disable learning, SE -stats enabled,
                         AD -packet action drop, LH -mac limit hit)

Logical          Index  Routing   Interface   Logical              MAC
interface               instance  device      Interface flags      limit
ge-11/0/3.0         79         3        136                          1024
ge-11/1/4.100       84         3        150                          1024
ge-11/1/1.100       86         3        147                          1024
ge-11/1/0.100       87         3        146                          1024
xe-10/2/0.100       88         3        144                          1024
xe-10/0/0.100       89         3        129                          1024
ge-11/1/0.200       90         4        146                          1024
ge-11/1/1.200       91         4        147                          1024
ge-11/1/4.200       92         4        150                          1024
xe-10/0/0.200       93         4        129                          1024
xe-10/2/0.200       94         4        144                          1024
```

show l2 learning-interface

```
user@host> run show l2-learning interface
Routing Instance Name : default-switch
Logical Interface flags (DL -disable learning, AD -packet action drop,
                        LH - MAC limit hit, DN - Interface Down, MP - MAC Pinning
enabled)
Logical            BD         MAC        STP          Logical
Interface          Name       Limit      State        Interface flags
ae0.0                         8192                      MP
```

show mac-rewrite interface

Syntax

show mac-rewrite interface
<brief | detail>
<*interface-name*>

Release Information

Command introduced in Junos OS Release 9.1.
Command introduced in Junos OS Release 14.1X53-D10 for EX4300 switches.
Command introducced in Junos OS Release 15.1X53-D55 for EX2300 and EX3400 switches.
Command introduced in Junos OS Release 17.4R1 for EX4600 switches.

Description

Display Layer 2 protocol tunneling (L2PT) information.

Options

brief | detail—(Optional) Display the specified level of output.

interface *interface-name*—(Optional) Display L2PT information for the specified interface.

Required Privilege Level

view

Related Documentation

- *layer2-control*
- *mac-rewrite*
- *protocol*
- *Understanding Layer 2 Protocol Tunneling Through a Network*
- *Understanding Layer 2 Protocol Tunneling Configuration Guidelines*
- *Configuring Layer 2 Protocol Tunneling*
- *Understanding Layer 2 Protocol Tunneling on EX Series Switches*
- *Configuring Layer 2 Protocol Tunneling on EX Series Switches with ELS Support (CLI Procedure)*

List of Sample Output

show mac-rewrite interface on page 2493
show mac-rewrite interface (EX Series Switches) on page 2493

Output Fields

Table 233 on page 2492 lists the output fields for the **show mac-rewrite interface** command. Output fields are listed in the approximate order in which they appear.

Table 233: show mac-rewrite interface Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface on which L2PT is configured.	**brief detail**

 Copyright © 2018, Juniper Networks, Inc.

Table 233: show mac-rewrite interface Output Fields (continued)

Field Name	Field Description	Level of Output
Protocols	Layer 2 protocols being tunneled on this interface. All devices that support L2PT can tunnel the following protocols: Cisco Discovery Protocol (CDP), Spanning Tree Protocol (STP), or VLAN Trunk Protocol (VTP). The following Layer 2 protocols can also be tunneled on some devices that support L2PT: E-LMI, GVRP, IEEE 802.1X, IEEE 802.3AH, LACP, LLDP, MMRP, MVRP, PVSTP+, UDLD, or VSTP. See *protocol* for more information on the supported protocols for tunneling on different devices.	**brief detail**

Sample Output

show mac-rewrite interface

```
user@host> show mac-rewrite interface
Interface       Protocols

 ge-1/0/5       STP VTP CDP PVSTP+
```

show mac-rewrite interface (EX Series Switches)

```
user@switch> show mac-rewrite interface
Interface       Protocols

 ge-0/0/1       802.3AH LLDP STP
```

Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

CHAPTER 30

VPN Operational Commands

Copyright © 2018, Juniper Networks, Inc.

clear vpls mac-address

Syntax

clear vpls mac-address
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<*mac-address*>

Release Information Command introduced before Junos OS Release 7.4.

Description (T Series and M Series routers, except for the M160 router) Clear media access control (MAC) address entries from the virtual private LAN service (VPLS) table.

Options

none—Clear all MAC address entries from the VPLS table for all routing instances.

instance *instance-name*—(Optional) Clear all MAC address entries for a VPLS instance from the VPLS table.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

mac-address—(Optional) Clear a specific MAC address in a VPLS instance from the VPLS table.

Required Privilege Level maintenance

List of Sample Output clear vpls mac-address on page 2496

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear vpls mac-address

```
user@host> clear vpls mac-address
```


Copyright © 2018, Juniper Networks, Inc.

clear vpls mac-table

Syntax

clear vpls mac-table
<instance *instance-name*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>
<*mac-address*>
<*vlan-id*>

Release Information

Command introduced before Junos OS Release 9.5.

Description

(MX Series routers) Clear media access control (MAC) addresses from the virtual private LAN service (VPLS) MAC table.

Options

none—Clear all MAC addresses from the VPLS table for all routing instances.

instance *instance-name*—(Optional) Clear all MAC addresses for a VPLS instance from the VPLS table.

interface *interface-name*—(Optional) Clear all MAC addresses for a VPLS interface from the VPLS table.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

mac-address—(Optional) Clear a specific MAC address in a VPLS instance from the VPLS table.

vlan-id—(Optional) Clear MAC addresses on a specified VLAN (0 through 4095).

Required Privilege Level

maintenance

List of Sample Output

clear vpls mac-table on page 2497

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear vpls mac-table

```
user@host> clear vpls mac-table
```

request l2circuit-switchover

Syntax

request l2circuit-switchover
<logical-system (all | logical-system-name)>
<neighbor *address*>
<virtual-circuit-id *identifier*>

Release Information

Command introduced in Junos OS Release 9.2.
Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.

Description

Manually trigger a switch from the active pseudowire to the redundant pseudowire. This command can be useful when performing network maintenance.

Options

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

neighbor *address*—(Optional) Trigger a switch of all of the active pseudowire connections with the specified neighbor to their respective redundant pseudowires.

virtual-circuit-id *identifier*—(Optional) Trigger a switch from the active pseudowire connection of the specified Layer 2 circuit to its redundant pseudowire.

Required Privilege Level

maintenance

Related Documentation

- *MPLS Feature Support on QFX Series and EX4600 Switches*

List of Sample Output

request l2circuit-switchover virtual-circuit-id on page 2498

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request l2circuit-switchover virtual-circuit-id

```
user@host> request l2circuit-switchover virtual-circuit-id 12
```


Copyright © 2018, Juniper Networks, Inc.

show dynamic-tunnels database

Syntax

show dynamic-tunnels database
<*destination*>
<logical-system (all | *logical-system-name*) >
<table *routing-table-name*>

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display dynamic tunnel database information.

Options

none—Display dynamic tunnel database information for all destinations and routing tables.

destination—(Optional) Display database entries for the specified IP address (with optional destination prefix length) only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

table *routing-table-name*—(Optional) Display database entries for the specified table only.

Required Privilege Level

view

List of Sample Output

show dynamic-tunnels database (Tunnel Is Up) on page 2500
show dynamic-tunnels database (No Tunnel PIC) on page 2500
show dynamic-tunnels database (Tunnel Is Expiring) on page 2501
show dynamic-tunnels database (Destination Specified) on page 2501
show dynamic-tunnels database (Localization) on page 2501

Output Fields

Table 234 on page 2499 lists the output fields for the **show dynamic-tunnels database** command. Output fields are listed in the approximate order in which they appear.

Table 234: show dynamic-tunnels database Output Fields

Field Name	Field Description
Table	Name of the routing table (for example, **inet.0**).
Destination-network	Destination IP address and subnet.
Tunnel to	Destination IP address and prefix of the tunnel.
State	State of the tunnel: **Up**, **Up (expires in *nn*:*nn*:*nn*seconds)**, or **Dn** (down).
Reference count	Number of routes across the dynamic tunnel that are currently being resolved.

Copyright © 2018, Juniper Networks, Inc.

Table 234: show dynamic-tunnels database Output Fields (continued)

Field Name	Field Description
Next-hop type	Type of tunnel: GRE or UDP (BGP-Signal).
Source address	Source IP address of the tunnel.
Next-hop	IP address of the destination interface.
VPN Label	The label provided by the peer device to identify the VPN through which the packet needs to go. This label is used to identify the VRF for route lookup.
Ingress Route	The IGP route along with the corresponding metric that has been selected for forwarding the tunnel-encapsulated packet.
Localized PFE	Packet Forwarding Engine interface which is the anchor Packet Forwarding Engine for the localized next-hop-based dynamic tunnels. When the anchor Packet Forwarding Engine of the tunnel goes down, it is represented by a **#** near the Packet Forwarding Engine name.
State	State of the destination interface: **Up**, **Dn**, or **Dn (no tunnel pic)**.

Sample Output

show dynamic-tunnels database (Tunnel Is Up)

```
user@host> show dynamic-tunnels database
Table: inet.3

Destination-network: 10.255.120.94/32
Tunnel to: 10.255.120.94/32
  Reference count: 4
  Next-hop type: UDP
    Source address: 10.255.120.92
    Next hop: tunnel-composite, 0x31132f64, nhid 3406
      VPN Label: Push 120 Reference count: 3
      Ingress Route: 10.255.120.94/32, via metric 2
      Traffic Statistics: Packets 241367951, Bytes 356741831578
      State: Up
```

show dynamic-tunnels database (No Tunnel PIC)

```
user@host> show dynamic-tunnels database
Table: inet.3

Destination-network: 10.255.120.94/32
Tunnel to: 10.255.120.94/32 State: Dn
  Reference count: 2
  Next-hop type: gre
    Source address: 10.255.120.92
    State: Dn (no tunnel pic)
```


Copyright © 2018, Juniper Networks, Inc.

show dynamic-tunnels database (Tunnel Is Expiring)

```
user@host> show dynamic-tunnels database
Table: inet.3

Destination-network: 10.255.120.94/32
Tunnel to: 10.255.120.94/32 State: Up (expires in 00:14:56 seconds)
  Reference count: 0
  Next-hop type: gre
    Source address: 10.255.120.92
    Next hop: gr-4/3/0.32769
    State: Up
```

show dynamic-tunnels database (Destination Specified)

```
user@host> show dynamic-tunnels database 10.255.120.94
Table: inet.3

Destination-network: 10.255.120.94/32
Tunnel to: 10.255.120.94/32 State: Up
  Reference count: 2
  Next-hop type: gre
    Source address: 10.255.120.92
    Next hop: gr-4/3/0.32769
    State: Up
```

show dynamic-tunnels database (Localization)

```
user@host> show dynamic-tunnels database
Destination-network: 1.0.0.0/8
Tunnel to: 1.1.1.6/32
  Reference count: 5
  Next-hop type: UDP
 Source address: 1.1.1.2
 Next hop: tunnel-composite, 0xc807930, nhid 1016
 Localized PFE: pfe-1/0/0
  VPN Label: Push 299808 Reference count: 4
 Ingress Route: 1.1.1.6/32, via metric 2
 Traffic Statistics: Packets 0, Bytes 0
 State: Up
```

show hfrr profiles

Syntax

show hfrr profiles
<brief |extensive>

Release Information

Command introduced in Junos OS Release 12.2.

Description

Display host fast reroute (HFRR) profile information.

HFRR adds a precomputed protection path into the Packet Forwarding Engine, such that if a link between a provider edge device and a server farm becomes unusable for forwarding, the Packet Forwarding Engine can use another path without having to wait for the router or the protocols to provide updated forwarding information.

Options

none—Display information about HRFF profiles.

brief | extensive—(Optional) Display the specified level of output.

Required Privilege Level

view

Related Documentation

- *Example: Configuring Link Protection with Host Fast Reroute*

List of Sample Output

show hfrr profiles on page 2503

Output Fields

Table 235 on page 2502 describes the output fields for the **show hfrr profiles** command. Output fields are listed in the approximate order in which they appear.

Table 235: show hfrr profiles Output Fields

Field Name	Field Description
HFRR	
HFRR current state	Status of the HFRR profile: HFRR_ACTIVE, HFRR_INACTIVE, HFRR_IFLH-NOT-CONF, and so on.
HFRR Prefix limit blackout timer expiry (in secs)	Time interval between an HFRR profile becoming inactive on exceeding the ARP prefix limit, and the profile starting the SYNC process.
HFRR prefix limit hit count	Number of times that an HFRR profile becomes inactive on exceeding the ARP prefix limit.
HFRR protected IFL name	Interface configured for the HFRR feature.
HFRR protected IFL handle	
HFRR routing instance name	The routing instance in which the HFRR interface is configured.

Copyright © 2018, Juniper Networks, Inc.

Table 235: show hfrr profiles Output Fields (continued)

Field Name	Field Description
HFRR routing instance handle	
HFRR sync BG scheduled	
HFRR RTS filter on	
HFRR delete BG scheduled	
HFRR ARP prefix limit	Configured ARP prefix limit.
HFRR ARP supplementary blackout timeout (in mins)	Supplementary time-out value configured for profile to be inactive when it hits ARP prefix limit.
HFRR number of ARP routes learned	Number of ARP routes learned on the configured interface.
HFRR number of FRR routes created	Number of ARP routes created on the configured interface.

Sample Output

show hfrr profiles

```
user@host> show hfrr profiles
HFRR pointer: 0x9254000
HFRR current state: HFRR_ACTIVE
HFRR Prefix limit blackout timer expiry (in secs): 0
HFRR prefix limit hit count: 0
HFRR protected IFL name: ge-4/1/0.0
HFRR protected IFL handle: 0x9248738
HFRR routing instance name: test
HFRR routing instance handle: 0x9145740
HFRR sync BG scheduled: NO
HFRR RTS filter on: YES
HFRR delete BG scheduled: NO
HFRR ARP prefix limit: 0
HFRR ARP supplementary blackout timeout (in mins): 1
HFRR number of ARP routes learned: 4
HFRR number of FRR routes created: 2
```

Copyright © 2018, Juniper Networks, Inc.

show ingress-replication mvpn

Syntax

show ingress-replication mvpn

Release Information

Command introduced in Junos OS Release 10.4.

Description

Display the state and configuration of the ingress replication tunnels created for the MVPN application when using the **mpls-internet-multicast** routing instance type.

Required Privilege Level

View

List of Sample Output

show ingress-replication mvpn on page 2504

Output Fields

Table 236 on page 2504 lists the output fields for the **show ingress-replication mvpn** command. Output fields are listed in the approximate order in which they appear.

Table 236: show ingress-replication mvpn Output Fields

Field Name	Field Description
Ingress tunnel	Identifies the MVPN ingress replication tunnel.
Application	Identifies the application (MVPN).
Unicast tunnels	List of unicast tunnels in use.
Leaf address	Address of the tunnel.
Tunnel type	Identifies the unicast tunnel type.
Mode	Indicates whether the tunnel was created as a new tunnel for the ingress replication, or if an existing tunnel was used.
State	Indicates whether the tunnel is Up or Down.

Sample Output

show ingress-replication mvpn

```
user@host> show ingress-replication mvpn
Ingress Tunnel: mvpn:1
  Application: MVPN
  Unicast tunnels
    Leaf Address       Tunnel-type        Mode        State
    10.255.245.2       P2P LSP            New         Up
    10.255.245.4       P2P LSP            New         Up
Ingress Tunnel: mvpn:2
  Application: MVPN
  Unicast tunnels
```


Copyright © 2018, Juniper Networks, Inc.

```
Leaf Address        Tunnel-type        Mode        State
10.255.245.2        P2P LSP            Existing    Up
```

show l2circuit connections

Syntax

show l2circuit connections
<brief | extensive | summary>
<down | up | up-down>
<history>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>
<neighbor *neighbor*>
<status>

Release Information

Command introduced before Junos OS Release 7.4.
Display enhancements in Junos OS Release 9.6.
Display enhancements in Junos OS Release 10.2.
Display enhancements in Junos OS Release 12.1.
Display enhancements in Junos OS Release 13.2.
Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.

Description

Display status information about Layer 2 virtual circuits from the local provider edge (PE) router to its neighbors.

Options

none—Display standard information about Layer 2 virtual circuits on all interfaces for all neighbors.

brief | extensive | summary—(Optional) Display the specified level of output. Use history to display information about connection history. Use status to display information about the connection and interface status.

down | up | up-down—(Optional) Display nonoperational, operational, or both kinds of connections.

history—(Optional) Display information about connection history.

interface *interface-name*—(Optional) Show all Layer 2 virtual circuits on an interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

neighbor *neighbor*—(Optional) IP address of a specific neighbor.

status—(Optional) Display information about the connection and interface status.

Required Privilege Level

view

List of Sample Output

show l2circuit connections on page 2510
show l2circuit connections interface on page 2511
show l2circuit connections extensive on page 2511

 Copyright © 2018, Juniper Networks, Inc.

show l2circuit connections extensive (Pseudowire Redundancy with Hot Standby) on page 2512

Output Fields Table 237 on page 2507 lists the output fields for the **show l2circuit connections** command. Output fields are listed in the approximate order in which they appear.

Table 237: show l2circuit connections Output Fields

Field Name	Field Description
Layer-2 Circuit Connections	Displays the legends for connection and interface status.
Neighbor	Remote PE neighbor.
Interface	Logical PE-to-CE interface on which the virtual circuit is configured.
Type	VC type: **rmt** (remote) or **loc** (local).

Copyright © 2018, Juniper Networks, Inc.

Table 237: show l2circuit connections Output Fields (continued)

Field Name	Field Description
Legend for connection status (St)	Status of the virtual circuit connection: • **EI**—The local virtual circuit interface is configured with an encapsulation that is not supported. • **MM**—The two routers do not agree on an MTU value, which causes an MTU mismatch. • **EM**—The encapsulation type received on this virtual circuit from the neighbor does not match the local virtual circuit interface encapsulation type. • **CM**—The two routers do not agree on a control word, which causes a control word mismatch. • **VM**—The remote and local VLAN IDs do not match across the Layer 2 circuit. • **OL**—No advertisement has been received for this virtual circuit from the neighbor. There is no outgoing label available for use by this virtual circuit. • **NC**—The interface is not configured as a CCC or TCC interface. • **BK**—The virtual circuit has switched to a backup connection. • **CB**—The remote PE router is advertising a different cell bundle from that configured on the local PE router. • **LD**—The connection to the local site is signaled down, because the CE-facing interface to the local site is down. • **RD**—The remote neighbor is down. It has signaled a problem using the pseudowire status code. • **NP**—The router detects that interface hardware is not present. The hardware may be offline, a PIC may not be of the desired type, or the interface may be configured in a different routing instance. • **Dn**—The virtual circuit is down. • **VC-Dn**—The virtual circuit is down because there is no tunnel LSP from the local PE router to the neighbor. • **UP**—The virtual circuit is operational. • **CF**—The router cannot find enough bandwidth to the remote router to satisfy the Layer 2 circuit bandwidth requirement. • **IB**—The bit rate is incompatible for Time Division Multiplexing (TDM). • **TDM**—TDM is not configured correctly. • **ST**—The virtual circuit has been switched to a standby connection. • **SP**—The virtual circuit connection is using a static pseudowire. • **RS**—The remote site is in a standby state. • **XX**—The virtual circuit is down for an unknown reason. This is a programming error.
Time last up	Date and time the virtual circuit was last operational.
# Up trans	Number of times the virtual circuit came up.
local-interface-name	Name of the local PE-to-CE interface.
Status	Status of the local interface.
Up	Interface is operational.

Copyright © 2018, Juniper Networks, Inc.

Table 237: show l2circuit connections Output Fields (continued)

Field Name	Field Description
Dn	Interface is not operational.
NP	Not present. Interface does not exist.
DS	Disabled. Interface has been administratively disabled.
WE	Wrong encapsulation. The interface is not configured as CCC.
UN	Interface status is initialized.
Encapsulation	Encapsulation of the local interface.
Flow Label Transmit	Flow label transmit status.
Flow Label Receive	Flow label receive status.
Remote PE	Prefix of the remote PE router.
Negotiated control-word	Whether the use of the control word has been negotiated for this virtual circuit: **Yes (Null)** or **No**.
Incoming label	Label used by the remote side of the virtual circuit to send packets destined to the local side. This label is routed to the local virtual circuit interface.
Outgoing label	Label used by the local side of the virtual circuit to send packets to the remote side of the virtual circuit. Packets originated on the local virtual circuit interface are encapsulated with this label before being placed on the tunnel LSP to the neighbor for this virtual circuit. This label is allocated by the neighbor and is used in demultiplexing incoming packets destined for this virtual circuit.
Negotiated PW status TLV	Displays the pseudowire status type, length, and value (TLV). TLVs are a method of encoding variable-length or optional information. If the pseudowire status TLV is used, the corresponding local or neighbor PE router status code is also displayed.
local PW status code	If the pseudowire status TLV is used, displays the local PE router status code.
Neighbor PW status code	If the pseudowire status TLV is used, displays the neighbor PE router status code.
Local interface	Name of the local interface used for the Layer 2 circuit connection.
Status	Status of the local interface (**Up** or **Down**).
Encapsulation	Encapsulation configured for the local interface.
APS-active	Indicates that the interface belongs to the working circuit.

Copyright © 2018, Juniper Networks, Inc.

Table 237: show l2circuit connections Output Fields (continued)

Field Name	Field Description
APS-inactive	Indicates that the interface belongs to the protect circuit.
Connection protection	Whether or not connection protection is configured for the Layer 2 circuit to the neighbor: **Yes** or **No**.
VC bandwidth	Bandwidth requirement of the Layer 2 circuit.
Time	Time at which the event occurred.
Connection History	Event types logged in history. • **loc intf up**—Local virtual circuit interface went up. • **loc intf down**—Local virtual circuit interface went down. • **In lbl Update**—Incoming label has been updated. • **Out lbl Update**—Outgoing label has been updated. • **PE route changed**—Route to PE router has been updated. • **PE route down**—Route to PE router is down. • **rmt side marked**—Remote side is marked. • **VC Dn**—Remote side indicated that its end of the virtual circuit is down (if the tunnel LSP from the remote side to the local side is down). • **status update timer**—Status update timer processing. It computes the state of the virtual circuit, and determines whether it should be advertised to or withdrawn from the remote side.

Sample Output

show l2circuit connections

```
user@host> show l2circuit connections
Layer-2 Circuit Connections:

Legend for connection status (St)
EI -- encapsulation invalid      NP -- interface h/w not present
MM -- mtu mismatch               Dn -- down
EM -- encapsulation mismatch     VC-Dn -- Virtual circuit Down
CM -- control-word mismatch      Up -- operational
VM -- vlan id mismatch           CF -- Call admission control failure
OL -- no outgoing label          IB -- TDM incompatible bitrate
NC -- intf encaps not CCC/TCC    TM -- TDM misconfiguration
BK -- Backup Connection          ST -- Standby Connection
CB -- rcvd cell-bundle size bad  SP -- Static Pseudowire
LD -- local site signaled down   RS -- remote site standby
RD -- remote site signaled down  HS -- hot standby
XX -- unknown

Legend for interface status
Up -- operational
Dn -- down
Neighbor: 10.255.245.51
    Interface                 Type  St     Time last up          # Up trans
    ge-2/0/2.600(vc 5)        rmt   Up     Dec  7 18:11:18 2009           1
```


Copyright © 2018, Juniper Networks, Inc.

```
        Remote PE: 10.255.245.51, Negotiated control-word: No
        Incoming label: 299856, Outgoing label: 299808
        Negotiated PW status TLV: No
        Local interface: ge-2/0/2.600, Status: Up, Encapsulation: VLAN
        Flow Label Transmit: No, Flow Label Receive: No
        Auto-sensed or Programmed by XYZ
```

Sample Output

show l2circuit connections interface

```
user@host> show l2circuit connections interface t1-2/0/0:1:1.0
Layer-2 Circuit Connections:

Legend for connection status (St)
EI -- encapsulation invalid      NP -- interface h/w not present
MM -- mtu mismatch               Dn -- down
EM -- encapsulation mismatch     VC-Dn -- Virtual circuit Down
CM -- control-word mismatch      Up -- operational
VM -- vlan id mismatch           CF -- Call admission control failure
OL -- no outgoing label          IB -- TDM incompatible bitrate
NC -- intf encaps not CCC/TCC    TM -- TDM misconfiguration
BK -- Backup Connection          ST -- Standby Connection
CB -- rcvd cell-bundle size bad  SP -- Static Pseudowire
LD -- local site signaled down   RS -- remote site standby
RD -- remote site signaled down  HS -- hot standby
XX -- unknown

Legend for interface status
Up -- operational
Dn -- down
Neighbor: 10.1.1.1
    Interface                 Type  St     Time last up          # Up trans
    t1-2/0/0:1:1.0(vc 1)(SP)  rmt   Up     Apr 27 04:21:02 2011           1
      Remote PE: 10.1.1.1, Negotiated control-word: Yes (Non-null)
      Incoming label: 1010001, Outgoing label: 1000001
      Negotiated PW status TLV: No
      Local interface: t1-1/0/0:1:1.0, Status: Up, Encapsulation: SATOP-T1,
APS-active
      Local interface: t1-2/0/0:1:1.0, Status: Up, Encapsulation: SATOP-T1,
APS-inactive
```

Sample Output

show l2circuit connections extensive

```
user@host>show l2circuit connections extensive
Layer-2 Circuit Connections:

Legend for connection status (St)
EI -- encapsulation invalid      NP -- interface h/w not present
MM -- mtu mismatch               Dn -- down
EM -- encapsulation mismatch     VC-Dn -- Virtual circuit Down
CM -- control-word mismatch      Up -- operational
VM -- vlan id mismatch           CF -- Call admission control failure
OL -- no outgoing label          IB -- TDM incompatible bitrate
NC -- intf encaps not CCC/TCC    TM -- TDM misconfiguration
BK -- Backup Connection          ST -- Standby Connection
CB -- rcvd cell-bundle size bad  SP -- Static Pseudowire
LD -- local site signaled down   RS -- remote site standby
```

```
RD -- remote site signaled down  HS -- hot standby
XX -- unknown

Legend for interface status
Up -- operational
Dn -- down
Neighbor: 10.255.49.149
    Interface                 Type  St     Time last up          # Up trans
    ae0.0(vc 100)             rmt   Up     Aug 31 09:36:12 2009           1
      Remote PE: 10.255.49.149, Negotiated control-word: Yes (Null)
      Incoming label: 299824, Outgoing label: 299776
      Negotiated PW status TLV: Yes
      local PW status code: 0x00000000, Neighbor PW status code: 0x00000000
      Local interface: ae0.0, Status: Up, Encapsulation: ETHERNET
      Connection protection: Yes
    Connection History:
        Aug 31 09:36:12 2009  status update timer
        Aug 31 09:36:12 2009  PE route changed
        Aug 31 09:36:12 2009  Out lbl Update                     299776
        Aug 31 09:36:12 2009  In lbl Update                      299824
        Aug 31 09:36:12 2009  loc intf up                         ae0.0
```

Sample Output

show l2circuit connections extensive (Pseudowire Redundancy with Hot Standby)

```
user@host>show l2circuit connections extensive
Layer-2 Circuit Connections:

Legend for connection status (St)
EI -- encapsulation invalid      NP -- interface h/w not present
MM -- mtu mismatch               Dn -- down
EM -- encapsulation mismatch     VC-Dn -- Virtual circuit Down
CM -- control-word mismatch      Up -- operational
VM -- vlan id mismatch           CF -- Call admission control failure
OL -- no outgoing label          IB -- TDM incompatible bitrate
NC -- intf encaps not CCC/TCC    TM -- TDM misconfiguration
BK -- Backup Connection          ST -- Standby Connection
CB -- rcvd cell-bundle size bad  SP -- Static Pseudowire
LD -- local site signaled down   RS -- remote site standby
RD -- remote site signaled down  HS -- Hot-standby Connection
XX -- unknown

Legend for interface status
Up -- operational
Dn -- down
Neighbor: 192.0.2.101
    Interface                 Type  St     Time last up          # Up trans
    ge-1/3/2.600(vc 1)        rmt   Up     Jan 24 11:00:26 2013           1
      Remote PE: 192.0.2.101, Negotiated control-word: Yes (Null)
      Incoming label: 299776, Outgoing label: 299776
      Negotiated PW status TLV: Yes
      local PW status code: 0x00000000, Neighbor PW status code: 0x00000000
      Local interface: ge-1/3/2.600, Status: Up, Encapsulation: VLAN
    Connection History:
        Jan 24 11:00:26 2013  status update timer
        Jan 24 11:00:26 2013  PE route changed
        Jan 24 11:00:26 2013  Out lbl Update                     299776
        Jan 24 11:00:26 2013  In lbl Update                      299776
        Jan 24 11:00:26 2013  loc intf up                  ge-1/3/2.600
Neighbor: 192.0.2.102
```


Copyright © 2018, Juniper Networks, Inc.

```
Interface                 Type  St     Time last up          # Up trans
ge-1/3/2.600(vc 2)        rmt   HS     -----                       ----
  Remote PE: 192.0.2.102, Negotiated control-word: Yes (Null)
  Incoming label: 299792, Outgoing label: 299776
  Negotiated PW status TLV: Yes
  local PW status code: 0x00000020, Neighbor PW status code: 0x00000000
  Local interface: ge-1/3/2.600, Status: Up, Encapsulation: VLAN
```

Copyright © 2018, Juniper Networks, Inc.

show l2vpn connections

Syntax

show l2vpn connections
<brief | extensive>
<down | up | up-down>
<history>
<instance *instance*>
<instance-history>
<local-site *local-site*>
<logical-system (all | *logical-system-name*)>
<remote-site *remote-site*>
<status>
<summary>

Release Information

Command introduced before Junos OS Release 7.4.
instance-history option introduced in Junos OS Release 12.3R2.

Description

Display Layer 2 virtual private network (VPN) connections.

Options

none—Display all Layer 2 VPN connections for all routing instances.

brief | extensive—(Optional) Display the specified level of output.

down | up | up-down—(Optional) Display nonoperational, operational, or both kinds of connections.

history—(Optional) Display information about connection history.

instance ***instance***—(Optional) Display connections for the specified routing instance only.

instance-history—(Optional) Display information about connection history for a particular instance.

local-site ***local-site***—(Optional) Display connections for the specified Layer 2 VPN local site name or ID only.

logical-system (all | ***logical-system-name*****)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

remote-site ***remote-site***—(Optional) Display connection for the specified Layer 2 VPN remote site ID only.

status—(Optional) Display information about the connection and interface status.

summary—(Optional) Display summary of all Layer 2 VPN connections information.

Required Privilege Level

view

List of Sample Output

show l2vpn connections on page 2517

Copyright © 2018, Juniper Networks, Inc.

show l2vpn connections on page 2518
show l2vpn connections extensive on page 2519
show l2vpn connections extensive (VPWS) on page 2519

Output Fields

Table 238 on page 2515 lists the output fields for the **show l2vpn connections** command. Output fields are listed in the approximate order in which they appear.

Table 238: show l2vpn connections Output Fields

Field Name	Field Description
Instance	Name of Layer 2 VPN instance.
L2vpn-id	For BGP autodiscovery, a globally unique Layer 2 VPN community identifier for the instance.
Local-ID	BGP **local-address** assigned to the local routing device.
Local site	Name of local site.
Local source-attachment-id	For FEC 129, the VPWS source attachment identifier. The point-to-point nature of VPWS requires that you specify the source access individual identifier (SAII) and the target access individual identifier (TAII). This SAII-TAII pair defines a unique pseudowire between two PE devices.
Target-attachment-id	For FEC 129, the VPWS target attachment identifier. If the configured target identifier matches a source identifier advertised by a remote PE device by way of a BGP auto-discovery message, the pseudowire between that source-target pair is signaled. If there is no match between an advertised source identifier and the configured target identifier, the pseudowire is not established.
Interface name	Name of interface.
Remote Site ID	Remote site ID.
Label Offset	Numbers within the label block that are skipped to find the next label base.
Label-base	Advertises the first label in a block of labels. A remote PE router uses this first label when sending traffic toward the advertising PE router.
Range	Advertises the label block size.
status-vector	Bit vector advertising the state of local PE-CE circuits to remote PE routers. A bit value of **0** indicates that the local circuit and LSP tunnel to the remote PE router are up, whereas a value of **1** indicates either one or both are down.
connection-site	Name of the connection site.
Type	Type of connection: **loc** (local) or **rmt** (remote).
St	Status of the connection. (For a list of possible values, see the **Legend for connection status (St)** field.)

Copyright © 2018, Juniper Networks, Inc.

Table 238: show l2vpn connections Output Fields (continued)

Field Name	Field Description
Time last up	Time that the connection was last in the **Up** condition.
# Up trans	Number of transitions from **Down** to **Up** condition.
Local circuit	Address and status of local circuit.
Remote circuit	Address and status of remote circuit.
St	Status of the Layer 2 VPN connection (corresponds with Legend for Connection Status): • **EI**—The local Layer 2 VPN interface is configured with an encapsulation that is not supported. • **EM**—The encapsulation type received on this Layer 2 VPN connection from the neighbor does not match the local Layer 2 VPN connection interface encapsulation type. • **VC-Dn**—The virtual circuit is currently down. • **CM**—The two routers do not agree on a control word, which causes a control word mismatch. • **CN**—The virtual circuit is not provisioned properly. • **OR**—The label associated with the virtual circuit is out of range. • **OL**—No advertisement has been received for this virtual circuit from the neighbor. There is no outgoing label available for use by this virtual circuit. • **LD**—All of the CE-facing interfaces to the local site are down. Therefore, the connection to the local site is signaled as down to the other PE routers. No pseudowires can be established. • **RD**—All the interfaces to the remote neighbor are down. Therefore, the remote site has been signaled as down to the other PE routers. No pseudowires can be established. • **LN**—The local site has lost path selection to the remote site and therefore no pseduowires can be established from this local site. • **RN**—The remote site has lost path selection to a local site or other remote site and therefore no pseudowires are established to this remote site. • **XX**—The Layer 2 VPN connection is down for an unn reason. This is a programming error. • **NC**—The interface encapsulation is not configured as an appropriate CCC, TCC, or Layer 2 VPN encapsulation. • **WE**—The encapsulation configured for the interface does not match the encapsulation configured for the associated connection within the Layer 2 VPN routing instance. • **NP**—The router detects that interface hardware is not present. The hardware might be offline, a PIC might not be of the desired type, or the interface might be configured in a different routing instance. • **->**—Only the outbound connection is up. • **<-**—Only the inbound connection is up. • **Up**—The Layer 2 VPN connection is operational. • **Dn**—The Layer 2 VPN connection is down. • **CF**—The router cannot find enough bandwidth to the remote router to satisfy the Layer 2 VPN connection bandwidth requirement.

Copyright © 2018, Juniper Networks, Inc.

Table 238: show l2vpn connections Output Fields (continued)

Field Name	Field Description
	• **SC**—The local site identifier matches the remote site identifier. No pseudowire can be established between these two sites. You should configure different values for the local and remote site identifiers. • **LM**—The local site identifier is not the minimum designated, meaning it is not the lowest. There is another local site with a lower site identifier. Pseudowires are not being established to this local site. and the associated local site identifier is not being used to distribute Layer 2 VPN label blocks. However, this is not an error state. Traffic continues to be forwarded to the PE router interfaces connected to the local sites when the local sites are in this state. • **RM**—The remote site identifier is not the minimum designated, meaning it is not the lowest. There is another remote site connected to the same PE router which has lower site identifier. The PE router cannot established a pseduowire to this remote site and the associated remote site identifier cannot be used to distribute VPLS label blocks. However, this is not an error state. Traffic can continue to be forwarded to the PE router interface connected to this remote site when the remote site is in this state. • **IL**—The incoming packets for the Layer 2 VPN connection have no MPLS label.
Remote PE	Address of the remote provider edge router.
Incoming label	Name of the incoming label.
Outgoing label	Name of the outgoing label.
Egress Protection	Whether the given PVC is protected by connection protection logic using egress protection for BGP signaled layer 2 services.
Flow Label Receive	Capability to pop the flow label in the receive direction to the remote provider edge (PE) router
Flow Label Transmit	Capability to push the flow label in the transmit direction to the provider edge (PE) router
Time	Date and time of Layer 2 VPN connection event.
Event	Type of event.
Interface/Lbl/PE	Interface, label, or PE router.

Sample Output

show l2vpn connections

```
user@host> show l2vpn connections
L2VPN Connections :
Instance : vpna
Edge protection: Not-Primary
Local site: 2 (ce-2)
offset: 1, range: 3, label-base: 32768
```

```
    connection-site           Type  St  Time last up          # Up trans
    3 (3)                     loc   Up  Jul 18 20:45:46 2001           1
      Local circuit: fe-0/0/0.1, Status: Up
      Remote circuit: fe-0/0/3.0, Status: Up
    1                         rmt   Up  Jul 18 21:47:25 2001           1
      Local circuit: fe-0/0/0.0, Status: Up
      Remote PE: 192.0.2.1
      Incoming label: 32768, Outgoing label: 32769
Local site: 3 (ce-3)
offset: 1, range: 2, label-base: 33792
    connection-site           Type  St  Time last up          # Up trans
    2 (ce-b)                  loc   Up  Jul 18 20:45:46 2001           1
      Local circuit: fe-0/0/0.1, Status: Up
      Remote circuit: fe-0/0/3.0, Status: Up
    1                         rmt   Up  Jul 18 21:47:25 2001           1
      Local circuit: fe-0/0/3.1, Status: Up
      Remote PE: 192.0.2.1
      Incoming label: 33792, Outgoing label: 32770
```

show l2vpn connections

```
user@host> show l2vpn connections

Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down    NP -- interface hardware not present
CM -- control-word mismatch      -> -- only outbound connection is up
CN -- circuit not provisioned    <- -- only inbound connection is up
OR -- out of range               Up -- operational
OL -- no outgoing label          Dn -- down
LD -- local site signaled down   CF -- call admission control failure
RD -- remote site signaled down  SC -- local and remote site ID collision
LN -- local site not designated  LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status  IL -- no incoming label
MM -- MTU mismatch               MI -- Mesh-Group ID not available
BK -- Backup connection          ST -- Standby connection
PF -- Profile parse failure      PB -- Profile busy
RS -- remote site standby        SN -- Static Neighbor
LB -- Local site not best-site   RB -- Remote site not best-site
VM -- VLAN ID mismatch

Legend for interface status
Up -- operational
Dn -- down

Instance: l2vpn-inst
Edge protection: Not-Primary
  Local site: pe2 (2)
    connection-site           Type  St     Time last up          # Up trans
    1                         rmt   Up     Jun 22 14:46:50 2015           1
      Remote PE: 10.255.255.1, Negotiated control-word: Yes (Null)
      Incoming label: 800002, Outgoing label: 800003
      Local interface: ge-0/0/1.300, Status: Up, Encapsulation: VLAN
      Flow Label Transmit: Yes, Flow Label Receive: Yes
```

 Copyright © 2018, Juniper Networks, Inc.

show l2vpn connections extensive

```
user@host> show l2vpn connections extensive
L2VPN Connections:
Instance: vpn-a
Edge protection: Not-Primary
Local site: ce-a (1)
  Interface name      Remote Site ID
  fe-0/0/0.0                     2
  Label Offset      Offset     Range
         32768           1         2
    connection-site          Type  St  Time last up          # Up trans
    2                        rmt   Up  Aug  3 00:08:14 2001           1
      Local circuit: fe-0/0/0.0, Status: Up
      Remote PE: 192.168.24.1
      Incoming label: 32769, Outgoing label: 32768
      Egress Protection: Yes
        Time                 Event            Interface/Lbl/PE
        Aug  3 00:08:14 2001  PE route up
        Aug  3 00:08:14 2001  Out lbl Update              32768
        Aug  3 00:08:14 2001  In lbl Update               32769
        Aug  3 00:08:14 2001  ckt0 up                fe-0/0/0.0
```

show l2vpn connections extensive (VPWS)

```
user@host> show l2vpn connections
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down    NP -- interface hardware not present
CM -- control-word mismatch      -> -- only outbound connection is up
CN -- circuit not provisioned    <- -- only inbound connection is up
OR -- out of range               Up -- operational
OL -- no outgoing label          Dn -- down
LD -- local site signaled down   CF -- call admission control failure
RD -- remote site signaled down  SC -- local and remote site ID collision
LN -- local site not designated  LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status  IL -- no incoming label
MM -- MTU mismatch               MI -- Mesh-Group ID not available
BK -- Backup connection          ST -- Standby connection
PF -- Profile parse failure      PB -- Profile busy
RS -- remote site standby        SN -- Static Neighbor
LB -- Local site not best-site   RB -- Remote site not best-site
VM -- VLAN ID mismatch

Legend for interface status
Up -- operational
Dn -- down

Instance: FEC129-VPWS
  L2vpn-id: 100:100
    Number of local interfaces: 1
    Number of local interfaces up: 1
    ge-2/0/5.0
  Local source-attachment-id: 1 (ONE)
    Target-attachment-id      Type  St     Time last up           # Up trans
    2                         rmt   Up     Nov 28 16:16:14 2012           1
```

```
  Remote PE: 198.51.100.2, Negotiated control-word: No
  Incoming label: 299792, Outgoing label: 299792
  Local interface: ge-2/0/5.0, Status: Up, Encapsulation: ETHERNET
Connection History:
    Nov 28 16:16:14 2012  status update timer
    Nov 28 16:16:14 2012  PE route changed
    Nov 28 16:16:14 2012  Out lbl Update                    299792
    Nov 28 16:16:14 2012  In lbl Update                     299792
    Nov 28 16:16:14 2012  loc intf up                   ge-2/0/5.0
```

 Copyright © 2018, Juniper Networks, Inc.

show mvpn c-multicast

Syntax

show mvpn c-multicast
<extensive | summary>
<instance-name *instance-name*>
<source-pe>

Release Information

Command introduced in Junos OS Release 8.4.
Option to show **source-pe** introduced in Junos OS Release 15.1.

Description

Display the multicast VPN customer multicast route information.

Options

extensive | summary—(Optional) Display the specified level of output.

instance-name *instance-name*—(Optional) Display output for the specified routing instance.

source-pe—(Optional) Display source-pe output for the specified c-multicast entries.

Required Privilege Level

view

List of Sample Output

show mvpn c-multicast on page 2522
show mvpn c-multicast summary on page 2522
show mvpn c-multicast extensive on page 2522
show mvpn c-multicast source-pe on page 2523

Output Fields

Table 239 on page 2521 lists the output fields for the **show mvpn c-multicast** command. Output fields are listed in the approximate order in which they appear.

Table 239: show mvpn c-multicast Output Fields

Field Name	Field Description	Level of Output
Instance	Name of the VPN routing instance.	summary extensive none
C-mcast IPv4 (S:G)	Customer router IPv4 multicast address.	extensive none
Ptnl	Provider tunnel attributes, *tunnel type:tunnel source, tunnel destination group*.	extensive none
St	State: • DS—Represents (S,G) and is created due to (*,G) • RM—Remote VPN route learned from the remote PE router • St display blank—SSM group join	extensive none
MVPN instance	Name of the multicast VPN routing instance	extensive none

Copyright © 2018, Juniper Networks, Inc.

Table 239: show mvpn c-multicast Output Fields (continued)

Field Name	Field Description	Level of Output
C-multicast IPv4 route count	Number of customer multicast IPv4 routes associated with the multicast VPN routing instance.	summary
C-multicast IPv6 route count	Number of customer multicast IPv6 routes associated with the multicast VPN routing instance.	summary

Sample Output

show mvpn c-multicast

```
user@host> show mvpn c-multicast
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
  C-mcast IPv4 (S:G)           Ptnl                           St
  192.168.195.78/32:203.0.113.1/24 PIM-SM:10.255.14.144, 198.51.100.1        RM
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
  C-mcast IPv4 (S:G)           Ptnl                           St
  192.168.195.94/32:203.0.113.0/24 PIM-SM:10.255.14.144, 198.51.100.2        RM
```

show mvpn c-multicast summary

```
user@host> show mvpn c-multicast summary
MVPN Summary:
Family: INET
Family: INET6

Instance: mvpn1
  C-multicast IPv6 route count: 1
```

show mvpn c-multicast extensive

```
user@host> show mvpn c-multicast extensive
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
  C-mcast IPv4 (S:G)           Ptnl                           St
```

 Copyright © 2018, Juniper Networks, Inc.

```
  192.168.195.78/32:203.0.113.1/24 PIM-SM:10.255.14.144, 198.51.100.1        RM
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
  C-mcast IPv4 (S:G)            Ptnl                           St
  192.168.195.94/32:203.0.113.0/24 PIM-SM:10.255.14.144, 198.51.100.2        RM
```

show mvpn c-multicast source-pe

```
user@host> show mvpn c-multicast source-pe
Family : INET
Family : INET6

Instance : mvpn1
  MVPN Mode : RPT-SPT
    C-Multicast route address: ::/0:ff05::1/128
        MVPN Source-PE1:
                extended-community: no-advertise target:10.1.0.0:9
                Route Distinguisher: 10.1.0.0:1
                Autonomous system number: 1
                Interface: ge-0/0/9.1 Index: 343
        PIM Source-PE1:
                extended-community: target:10.1.0.0:9
                Route Distinguisher: 10.1.0.0:1
                Autonomous system number: 1
                Interface: ge-0/0/9.1 Index: 343
```

Copyright © 2018, Juniper Networks, Inc.

show mvpn instance

Syntax

show mvpn instance
<*instance-name*>
<display-tunnel-name>
<extensive | summary>
<inet | inet6>
<logical-system>

Release Information

Command introduced in Junos OS Release 8.4.
Additional details in output for extensive option introduced in Junos OS Release 15.1.

Description

Display the multicast VPN routing instance information according the options specified.

Options

instance-name—(Optional) Display statistics for the specified routing instance, or press Enter without specifying an instance name to show output for all instances.

display-tunnel-name—(Optional) Display the ingress provider tunnel name rather than the attribute.

extensive | summary—(Optional) Display the specified level of output.

inet | inet6—(Optional) Display output for the specified IP type.

inet | inet6—(Optional) Display output for the specified IP type.

logical-system—(Optional) Display details for the specified logical system, or type “all”.

Required Privilege Level

view

List of Sample Output

show mvpn instance on page 2525
show mvpn instance summary on page 2526
show mvpn instance extensive on page 2526
show mvpn instance summary (IPv6) on page 2527

Output Fields

Table 240 on page 2524 lists the output fields for the **show mvpn instance** command. Output fields are listed in the approximate order in which they appear.

Table 240: show mvpn instance Output Fields

Field Name	Field Description	Level of Output
MVPN instance	Name of the multicast VPN routing instance	extensive none
Instance	Name of the VPN routing instance.	summary extensive none
Provider tunnel	Provider tunnel attributes, *tunnel type:tunnel source, tunnel destination group*.	extensive none

Copyright © 2018, Juniper Networks, Inc.

Table 240: show mvpn instance Output Fields (continued)

Field Name	Field Description	Level of Output
Neighbor	Address, type of provider tunnel (I-P-tnl, inclusive provider tunnel and S-P-tnl, selective provider tunnel) and provider tunnel for each neighbor.	extensive none
C-mcast IPv4 (S:G)	Customer IPv4 router multicast address.	extensive none
C-mcast IPv6 (S:G)	Customer IPv6 router multicast address.	extensive none
Ptnl	Provider tunnel attributes, *tunnel type:tunnel source, tunnel destination group*.	extensive none
St	State: • DS—Represents (S,G) and is created due to (*,G) • RM—Remote VPN route learned from the remote PE router • St display blank—SSM group join	extensive none
Neighbor count	Number of neighbors associated with the multicast VPN routing instance.	summary
C-multicast IPv4 route count	Number of customer multicast IPv4 routes associated with the multicast VPN routing instance.	summary
C-multicast IPv6 route count	Number of customer multicast IPv6 routes associated with the multicast VPN routing instance.	summary

Sample Output

show mvpn instance

```
user@host> show mvpn instance
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
  Provider tunnel: I-P-tnl:PIM-SM:10.255.14.144, 198.51.100.1
  Neighbor                              I-P-tnl
  10.255.14.160                         PIM-SM:10.255.14.160, 198.51.100.1
  10.255.70.17                          PIM-SM:10.255.70.17, 198.51.100.1
  C-mcast IPv4 (S:G)            Ptnl                          St
  192.168.195.78/32:203.0.113.0/24 PIM-SM:10.255.14.144, 198.51.100.1        RM
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
```

```
  Provider tunnel: I-P-tnl:PIM-SM:10.255.14.144, 198.51.100.2
  Neighbor                              I-P-tnl
  10.255.14.160                         PIM-SM:10.255.14.160, 198.51.100.2
  10.255.70.17                          PIM-SM:10.255.70.17, 198.51.100.2
  C-mcast IPv4 (S:G)           Ptnl                           St
  192.168.195.94/32:203.0.113.1/24 PIM-SM:10.255.14.144, 198.51.100.2        RM
```

Sample Output

show mvpn instance summary

```
user@host> show mvpn instance summary
MVPN Summary:
Family: INET
Family: INET6

Instance: mvpn1
  Sender-Based RPF: Disabled. Reason: Not enabled by configuration.
  Hot Root Standby: Disabled. Reason: Not enabled by configuration.
  Neighbor count: 3
  C-multicast IPv6 route count: 1
```

Sample Output

show mvpn instance extensive

```
user@host> show mvpn instance extensive
MVPN instance:
Family : INET

Instance : vpn_blue
    Customer Source: 10.1.1.1
        RT-Import Target: 192.168.1.1:100
        Route-Distinguisher: 192.168.1.1:100
        Source-AS: 65000
        Via unicast route: 10.1.0.0/16 in vpn-blue.inet.0
        Candidate Source PE Set:
            RT-Import 192.168.1.1:100, RD 1111:22222, Source-AS 65000
            RT-Import 192.168.2.2:100, RD 1111:22222, Source-AS 65000
            RT-Import 192.168.3.3:100, RD 1111:22222, Source-AS 65000

'Extensive' output will show everything in 'detail' output and add the list of
bound c-multicast routes.

> show mvpn source 10.1.1.1 instance vpn_blue extensive

Family : INET

Instance : vpn_blue
    Customer Source: 10.1.1.1
        RT-Import Target: 192.168.1.1:100
        Route-Distinguisher: 192.168.1.1:100
        Source-AS: 65000
        Via unicast route: 10.1.0.0/16 in vpn-blue.inet.0
        Candidate Source PE Set:
            RT-Import 192.168.1.1:100, RD 1111:22222, Source-AS 65000
            RT-Import 192.168.2.2:100, RD 1111:22222, Source-AS 65000
            RT-Import 192.168.3.3:100, RD 1111:22222, Source-AS 65000
        Customer-Multicast Routes:
```

 Copyright © 2018, Juniper Networks, Inc.

```
            10.1.1.1/32:198.51.100.3/24
            10.1.1.1/32:198.51.100.3/24
```

show mvpn instance summary (IPv6)

```
user@host> show mvpn instance summary
MVPN Summary:
Instance: VPN-A
  C-multicast IPv6 route count: 2
Instance: VPN-B
  C-multicast IPv6 route count: 2
```

show mvpn neighbor

Syntax

show mvpn neighbor
<extensive | summary>
<inet | inet6>
<instance *instance-name* | neighbor-address *address*>
<logical-system *logical-system-name*>

Release Information Command introduced in Junos OS Release 8.4.

Description Display multicast VPN neighbor information.

Options

extensive | summary—(Optional) Display the specified level of output for all multicast VPN neighbors.

inet | inet6—(Optional) Display IPv4 or IPv6 information for all multicast VPN neighbors.

instance *instance-name* | neighbor-address *address*—(Optional) Display multicast VPN neighbor information for the specified instance or the specified neighbor.

logical-system *logical-system-name*—(Optional) Display multicast VPN neighbor information for the specified logical system.

Required Privilege Level view

List of Sample Output

show mvpn neighbor on page 2529
show mvpn neighbor extensive on page 2529
show mvpn neighbor extensive on page 2530
show mvpn neighbor instance-name on page 2530
show mvpn neighbor neighbor-address on page 2530
show mvpn neighbor neighbor-address summary on page 2531
show mvpn neighbor neighbor-address extensive on page 2531
show mvpn neighbor neighbor-address instance-name on page 2531
show mvpn neighbor summary on page 2532

Output Fields Table 241 on page 2528 lists the output fields for the **show mvpn neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 241: show mvpn neighbor Output Fields

Field Name	Field Description	Level of Output
MVPN instance	Name of the multicast VPN routing instance	**extensive** none
Instance	Name of the VPN routing instance.	**summary extensive** none

Copyright © 2018, Juniper Networks, Inc.

Table 241: show mvpn neighbor Output Fields (continued)

Field Name	Field Description	Level of Output
Neighbor	Address, type of provider tunnel (**I-P-tnl**, inclusive provider tunnel and **S-P-tnl**, selective provider tunnel) and provider tunnel for each neighbor.	**extensive** none
Provider tunnel	Provider tunnel attributes, ***tunnel type:tunnel source, tunnel destination group***.	**extensive** none

Sample Output

show mvpn neighbor

```
user@host> show mvpn neighbor
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
  Neighbor                              I-P-tnl
  10.255.14.160                         PIM-SM:10.255.14.160, 192.0.2.1
  10.255.70.17                          PIM-SM:10.255.70.17, 192.0.2.1
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
  Neighbor                              I-P-tnl
  10.255.14.160                         PIM-SM:10.255.14.160, 192.0.2.2
  10.255.70.17                          PIM-SM:10.255.70.17, 192.0.2.2
```

Sample Output

show mvpn neighbor extensive

```
user@host> show mvpn neighbor extensive
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
  Neighbor                              I-P-tnl
  10.255.14.160                         PIM-SM:10.255.14.160, 192.0.2.1
  10.255.70.17                          PIM-SM:10.255.70.17, 192.0.2.1
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel
```

```
Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
  Neighbor                              I-P-tnl
  10.255.14.160                         PIM-SM:10.255.14.160, 192.0.2.2
  10.255.70.17                          PIM-SM:10.255.70.17, 192.0.2.2
```

show mvpn neighbor extensive

```
user@host> show mvpn neighbor extensive
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: mvpn-a
  Neighbor                              I-P-tnl
  10.255.72.45
  10.255.72.50                          LDP P2MP:10.255.72.50, lsp-id 1
```

Sample Output

show mvpn neighbor instance-name

```
user@host> show mvpn neighbor instance-name VPN-A
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
  Neighbor                              I-P-tnl
  10.255.14.160                         PIM-SM:10.255.14.160, 192.0.2.1
  10.255.70.17                          PIM-SM:10.255.70.17, 192.0.2.1
```

Sample Output

show mvpn neighbor neighbor-address

```
user@host> show mvpn neighbor neighbor-address 10.255.14.160
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
  Neighbor                              I-P-tnl
  10.255.14.160                         PIM-SM:10.255.14.160, 192.0.2.1
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel
```

 Copyright © 2018, Juniper Networks, Inc.

```
Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
  Neighbor                             I-P-tnl
  10.255.14.160                        PIM-SM:10.255.14.160, 192.0.2.2
```

Sample Output

show mvpn neighbor neighbor-address summary

```
user@host> show mvpn neighbor neighbor-address 10.255.70.17 summary
MVPN Summary:
Instance: VPN-A
Instance: VPN-B
```

Sample Output

show mvpn neighbor neighbor-address extensive

```
user@host> show mvpn neighbor neighbor-address 10.255.70.17 extensive
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
  Neighbor                             I-P-tnl
  10.255.70.17                         PIM-SM:10.255.70.17, 192.0.2.1
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
  Neighbor                             I-P-tnl
  10.255.70.17                         PIM-SM:10.255.70.17, 192.0.2.2
```

Sample Output

show mvpn neighbor neighbor-address instance-name

```
user@host> show mvpn neighbor neighbor-address 10.255.70.17 instance-name VPN-A
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
  Neighbor                             I-P-tnl
  10.255.70.17                         PIM-SM:10.255.70.17, 192.0.2.1
```

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show mvpn neighbor summary

```
user@host> show mvpn neighbor summary
MVPN Summary:
Family: INET
Family: INET6

Instance: mvpn1
  Neighbor count: 3
```

 Copyright © 2018, Juniper Networks, Inc.

show vpls connections

Syntax

show vpls connections
<brief | extensive>
<down | up | up-down>
<history>
<instance *instance-name* local-site *local-site-name* remote-site *remote-site-name*>
<instance-history>
<logical-system (all | *logical-system-name*)>
<status>
<summary>

Release Information

Command introduced before Junos OS Release 7.4.
instance-history option introduced in Junos OS Release 12.3R2.

Description

(T Series and M Series routers, except for the M160 router) Display virtual private LAN service (VPLS) connection information.

Options

none—Display information about all VPLS connections for all routing instances.

brief | extensive—(Optional) Display the specified level of output.

down | up | up-down—(Optional) Display nonoperational, operational, or both types of connections.

history—(Optional) Display information about connection history.

instance *instance-name*—(Optional) Display the VPLS connections for the specified routing instance only.

instance-history—(Optional) Display information about connection history for a particular instance.

local-site *local-site-name*—(Optional) Display the VPLS connections for the specified local site name or ID only.

remote-site *remote-site-name*—(Optional) Display the VPLS connections for the specified remote site name or ID only. Label block size information is always shown as 0 when using this option.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

status—(Optional) Display information about the connection and interface status.

summary—(Optional) Display summary of all VPLS connections information.

Required Privilege Level

view

Copyright © 2018, Juniper Networks, Inc.

List of Sample Output

show vpls connections on page 2538
show vpls connections (with FEC128 and FEC129 in the same routing-instance) on page 2540
show vpls connections (with multiple pseudowires) on page 2541
show vpls connections extensive (Static VPLS Neighbors) on page 2542

Output Fields

Table 242 on page 2534 lists the output fields for the **show vpls connections** command. Output fields are listed in the approximate order in which they appear.

Table 242: show vpls connections Output Fields

Field Name	Field Description
Instance	Name of the VPLS instance.
Local site	Name of the local site.
VPLS-id	Identifier for the VPLS site.
Number of local interfaces	Number of interfaces configured for the local site.
Number of local interfaces up	Number of interfaces configured for the local site that are currently up.
IRB interface present	Indicates whether or not an integrated routing and bridging (IRB) interface is present (**yes** or **no**).
Intf	List of all of the interfaces configured for the local site. The types of interfaces can include VPLS virtual loopback tunnel interfaces and label-switched interfaces. Any interface that supports VPLS could be listed here. Virtual loopback tunnel interfaces are displayed using the **vt-*fpc/pic/port.nnnnn*** format. Label-switched interfaces are displayed using the **lsi.*nnnnn*** format. In both cases, ***nnnnn*** is a dynamically generated virtual port used to transport and receive packets from other provider edge (PE) routers in the VPLS domain. Each interface might include the following information: • Identification as a VPLS interface • Name of the associated VPLS routing instance • Local site number • Remote site number • VPLS neighbor address • VPLS identifier
Interface flags	Flag associated with the interface. Can include the following: • **VC-Down**—The virtual circuit associated with this interface is down.
Label-base	First label in a block of labels. A remote PE router uses this first label when sending traffic toward the advertising PE router.

 Copyright © 2018, Juniper Networks, Inc.

Table 242: show vpls connections Output Fields (continued)

Field Name	Field Description
Offset	Displays the VPLS Edge (VE) block offset in the Layer 2 VPN NLRI. The VE block offset is used to identify a label block from which a particular label value is selected to setup a pseduowire for a remote site. The block offset value itself indicates the starting VE ID that maps to the label base contained in the VPLS NLRI advertisement.
Size	Label block size. A configurable value that represents the number of label blocks required to cover all the pseudowires for the remote peer. Acceptable configuration values are: **2**, **4**, **8** and **16**. The default value is **2**. A value of **0** will be displayed when using the **remote-site** option.
Range	Label block range. A value that keeps track of the numbers of remote sites discovered within each label block.
Preference	Preference value advertised for a VPLS site. When multiple PE routers are assigned the same VE ID for multihoming, you might need to specify that a particular PE router acts as the designated forwarder by configuring the site preference value. The site preference indicates the degree of preference for a particular customer site. The site preference is one of the tie-breaking criteria used in a designated forwarder election.
status-vector	Bit vector advertising the state of local PE-CE circuits to remote PE routers. A bit value of **0** indicates that the local circuit and LSP tunnel to the remote PE router are up, whereas a value of **1** indicates either one or both are down.
connection-site	Name of the connection site.
Neighbor	IP address and VPLS identifier for the VPLS neighbor. If multiple pseudowires have been configured, the IP address will also show the PW-specific *vpls-id-list*, for example, 203.0.113.144 (vpls-id 200).
Type	Type of connection: **loc** (local) or **rmt** (remote).

Copyright © 2018, Juniper Networks, Inc.

Table 242: show vpls connections Output Fields (continued)

Field Name	Field Description

St	Status of the VPLS connection (corresponds with Legend for Connection Status):
	
	• **EI**—The local VPLS interface is configured with an encapsulation that is not supported.
	• **EM**—The encapsulation type received on this VPLS connection from the neighbor does not match the local VPLS connection interface encapsulation type.
	• **VC-Dn**—The virtual circuit is currently down.
	• **CM**—The two routers do not agree on a control word, which causes a control word mismatch.
	• **CN**—The virtual circuit is not provisioned properly.
	• **OR**—The label associated with the virtual circuit is out of range.
	• **OL**—No advertisement has been received for this virtual circuit from the neighbor. There is no outgoing label available for use by this virtual circuit.
	• **LD**—All of the CE-facing interfaces to the local site are down. Therefore, the connection to the local site is signaled as down to the other PE routers. No pseudowires can be established.
	• **RD**—All the interfaces to the remote neighbor are down. Therefore, the remote site has been signaled as down to the other PE routers. No pseudowires can be established.
	• **LN**—The local site has lost path selection to the remote site and therefore no pseduowires can be established from this local site.
	• **RN**—The remote site has lost path selection to a local site or other remote site and therefore no pseudowires are established to this remote site.
	In a multihoming configuration, one multihomed PE site displays the state **LN**, and the other multihomed PE site displays the state **RN** in the following circumstances:
	  ◦ The multihomed links are both configured to be the backup site.
	  ◦ The two multihomed PE routers have the same site ID, but have a peering relationship with a route reflector (RR) that has a different site ID.
	• **XX**—The VPLS connection is down for an unn reason. This is a programming error.
	• **MM**—The MTU for the local site and the remote site do not match.
	• **BK**—The router is using a backup connection.
	• **PF**—Profile parse failure.
	• **RS**—The remote site is in a standby state.
	• **NC**—The interface encapsulation is not configured as an appropriate CCC, TCC, or VPLS encapsulation.
	• **WE**—The encapsulation configured for the interface does not match the encapsulation configured for the associated connection within the VPLS routing instance.

 Copyright © 2018, Juniper Networks, Inc.

Table 242: show vpls connections Output Fields (continued)

Field Name	Field Description
	• **NP**—The router detects that interface hardware is not present. The hardware might be offline, a PIC might not be of the desired type, or the interface might be configured in a different routing instance. • **-->**—Only the outbound connection is up. • **<-**—Only the inbound connection is up. • **Up**—The VPLS connection is operational. • **Dn**—The VPLS connection is down. • **CF**—The router cannot find enough bandwidth to the remote router to satisfy the VPLS connection bandwidth requirement. • **SC**—The local site identifier matches the remote site identifier. No pseudowire can be established between these two sites. You should configure different values for the local and remote site identifiers. • **LM**—The local site identifier is not the minimum designated, meaning it is not the lowest. There is another local site with a lower site identifier. Pseudowires are not being established to this local site. and the associated local site identifier is not being used to distribute VPLS label blocks. However, this is not an error state. Traffic continues to be forwarded to the PE router interfaces connected to the local sites when the local sites are in this state. • **RM**—The remote site identifier is not the minimum designated, meaning it is not the lowest. There is another remote site connected to the same PE router which has lower site identifier. The PE router cannot established a pseudowire to this remote site and the associated remote site identifier cannot be used to distribute VPLS label blocks. However, this is not an error state. Traffic can continue to be forwarded to the PE router interface connected to this remote site when the remote site is in this state. • **IL**—The incoming packets for the VPLS connection have no MPLS label. • **MI**—The configured mesh group identifier is in use by another system in the network. • **ST**—The router has switched to a standby connection. • **PB**—Profile busy. • **SN**—The VPLS neighbor is static.
Time last up	Time connection was last in the **Up** condition.
# Up trans	Number of transitions from **Down** to **Up** condition.
Status	Status of the (local or remote circuit) local interface: • **Up**—Operational • **Dn**—Down • **NP**—Not present • **DS**—Disabled • **WE**—Wrong encapsulation • **UN**—Uninitialized
Encapsulation	Type of encapsulation: **VPLS**.
Remote PE	Address of the remote provider edge router.

Table 242: show vpls connections Output Fields (continued)

Field Name	Field Description
Negotiated control-word	Whether a control word has been negotiated: **Yes** or **No**.
Incoming label	Name of the incoming label.
Outgoing label	Name of the outgoing label.
Negotiated PW status TLV	Indicates whether or not the pseudowire status TLV has been negotiated for the VPLS connection.
Local interface	Provides the following information about the local interface configured for the VPLS neighbor: • Name of the local interface • **Status**—Interface status (**Up** or **Down**) • **Encapsulation**—Interface encapsulation (for example, **ETHERNET**) • **Description**—Includes the VPLS instance name, the VPLS neighbor address, and the VPLS identifier
Time	Date and time of VPLS connection event.
Event	Type of event.
Interface/Lbl/PE	Interface, label, or PE router.
Connection History	Each entry can include the date, time, year, and the connection event. Connection events include any of a variety of events related to VPLS connections, such as route changes, label updates, and interfaces going down or coming up.

Sample Output

show vpls connections

```
user@host> show vpls connections
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down    NP -- interface hardware not present
CM -- control-word mismatch      -< -- only outbound connection is up
CN -- circuit not provisioned    >- -- only inbound connection is up
OR -- out of range               Up -- operational
OL -- no outgoing label          Dn -- down
LD -- local site signaled down   CF -- call admission control failure
RD -- remote site signaled down  SC -- local and remote site ID collision
LN -- local site not designated  LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unn connection status  IL -- no incoming label
MM -- MTU mismatch               MI -- Mesh-Group ID not available
BK -- Backup connection          ST -- Standby connection
```


Copyright © 2018, Juniper Networks, Inc.

```
PF -- Profile parse failure      PB -- Profile busy

Legend for interface status
Up -- operational
Dn -- down

Instance: vpls-1
  Local site: 1 (11)
    Number of local interfaces: 1
    Number of local interfaces up: 1
    IRB interface present: no
    lt-1/3/0.10496
    vt-1/3/0.1048588    1         Intf - vpls vpls-1 local site 11 remote site 1

    vt-1/2/0.1048591    2         Intf - vpls vpls-1 local site 11 remote site 2

    vt-1/2/0.1048585    3         Intf - vpls vpls-1 local site 11 remote site 3

    vt-1/2/0.1048587    4         Intf - vpls vpls-1 local site 11 remote site 4

    vt-1/2/0.1048589    5         Intf - vpls vpls-1 local site 11 remote site 5

    vt-1/3/0.1048586    6         Intf - vpls vpls-1 local site 11 remote site 6

    vt-1/3/0.1048590    7         Intf - vpls vpls-1 local site 11 remote site 7

    vt-1/3/0.1048584    8         Intf - vpls vpls-1 local site 11 remote site 8

    Label-base        Offset     Size  Range     Preference
    800256            1          16     16        100
    Timer Values:
        Startup wait time: 120 seconds
        New site wait-time: 20 seconds
        Collision detect time: 30 seconds
        Reclaim wait time: 748 milliseconds
    connection-site           Type  St     Time last up          # Up trans
    1                         rmt   Up     Apr 28 13:28:24 2009           2
      Remote PE: 192.0.2.1, Negotiated control-word: No
      Incoming label: 800256, Outgoing label: 800026
      Local interface: vt-1/3/0.1048588, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls vpls-1 local site 11 remote site 1
    Connection History:
        Apr 28 13:28:24 2009  status update timer
        Apr 28 13:28:24 2009  PE route down
        Apr 28 13:24:27 2009  status update timer
        Apr 28 13:24:27 2009  loc intf up            vt-1/3/0.1048588
        Apr 28 13:24:27 2009  PE route changed
        Apr 28 13:24:27 2009  Out lbl Update                    800026
        Apr 28 13:24:27 2009  In lbl Update                     800256
        Apr 28 13:24:27 2009  loc intf down
    2                         rmt   Up     Apr 28 13:28:24 2009           2
      Remote PE: 192.0.2.71, Negotiated control-word: No
      Incoming label: 800257, Outgoing label: 800034
      Local interface: vt-1/2/0.1048591, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls vpls-1 local site 11 remote site 2
    Connection History:
        Apr 28 13:28:24 2009  status update timer
        Apr 28 13:28:24 2009  PE route down
        Apr 28 13:24:28 2009  status update timer
        Apr 28 13:24:28 2009  loc intf up            vt-1/2/0.1048591
        Apr 28 13:24:28 2009  PE route changed
```

Copyright © 2018, Juniper Networks, Inc.

```
Apr 28 13:24:28 2009  Out lbl Update                    800034
Apr 28 13:24:28 2009  In lbl Update                     800257
Apr 28 13:24:28 2009  loc intf down
```

show vpls connections (with FEC128 and FEC129 in the same routing-instance)

```
user@host> show vpls connections
Instance: fec129
  L2vpn-id: 1:1
  Local-id: 203.0.113.0
FEC129-VPLS State:
  Mesh-group connections: __ves__
    Remote-id                 Type  St     Time last up          # Up trans
    203.0.3.3                  rmt   Up     Sep 19 09:59:56 2017           1
      Remote PE: 203.0.3.3, Negotiated control-word: No
      Incoming label: 262155, Outgoing label: 262164
      Negotiated PW status TLV: No
      Local interface: lsi.1048844, Status: Up, Encapsulation: ETHERNET
        Description: Intf - vpls fec129 local-id 10.4.4.4 remote-id 203.0.3.3
neighbor 203.0.3.3
      Flow Label Transmit: No, Flow Label Receive: No
    203.0.2.2                  rmt   Up     Sep 19 09:59:52 2017           1
      Remote PE: 203.0.2.2, Negotiated control-word: No
      Incoming label: 262154, Outgoing label: 262157
      Negotiated PW status TLV: No
      Local interface: lsi.1048846, Status: Up, Encapsulation: ETHERNET
        Description: Intf - vpls fec129 local-id 10.4.4.4 remote-id 203.0.2.2
neighbor 203.0.2.2
      Flow Label Transmit: No, Flow Label Receive: No
    203.0.1.1                  rmt   Up     Sep 19 09:59:48 2017           1
      Remote PE: 203.0.1.1, Negotiated control-word: No
      Incoming label: 262156, Outgoing label: 262157
      Negotiated PW status TLV: No
      Local interface: lsi.1048845, Status: Up, Encapsulation: ETHERNET
        Description: Intf - vpls fec129 local-id 10.4.4.4 remote-id 203.0.1.1
neighbor 203.0.1.1
      Flow Label Transmit: No, Flow Label Receive: No

LDP-VPLS State
   Mesh-group connections: MG1
     Neighbor                  Type  St     Time last up          # Up trans
     203.0.6.6(vpls-id 1)       rmt   Up     Sep 17 19:17:11 2017           1
       Remote PE: 203.0.6.6, Negotiated control-word: No
       Incoming label: 262423, Outgoing label: 262145
       Negotiated PW status TLV: No
       Local interface: lsi.1049859, Status: Up, Encapsulation: ETHERNET
         Description: Intf - vpls bgp-vpls neighbor 203.0.6.6 vpls-id 1
       Flow Label Transmit: No, Flow Label Receive: No
     203.0.7.7(vpls-id 1)       rmt   Up     Sep 17 19:17:04 2017           1
       Remote PE: 203.0.7.7, Negotiated control-word: No
       Incoming label: 262424, Outgoing label: 262145
       Negotiated PW status TLV: No
       Local interface: lsi.1049857, Status: Up, Encapsulation: ETHERNET
         Description: Intf - vpls bgp-vpls neighbor 203.0.7.7 vpls-id 1
       Flow Label Transmit: No, Flow Label Receive: No
   Mesh-group connections: MG2
     Neighbor                  Type  St     Time last up          # Up trans
     203.0.5.5(vpls-id 1)       rmt   Up     Sep 17 19:17:00 2017           1
       Remote PE: 203.0.5.5, Negotiated control-word: No
       Incoming label: 262425, Outgoing label: 299872
       Negotiated PW status TLV: No
```

 Copyright © 2018, Juniper Networks, Inc.

```
       Local interface: lsi.1049856, Status: Up, Encapsulation: ETHERNET
         Description: Intf - vpls bgp-vpls neighbor 203.0.5.5 vpls-id 1
       Flow Label Transmit: No, Flow Label Receive: No
```

show vpls connections (with multiple pseudowires)

```
user@host> show vpls connections
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down    NP -- interface hardware not present
CM -- control-word mismatch      -> -- only outbound connection is up
CN -- circuit not provisioned    <- -- only inbound connection is up
OR -- out of range               Up -- operational
OL -- no outgoing label          Dn -- down
LD -- local site signaled down   CF -- call admission control failure
RD -- remote site signaled down  SC -- local and remote site ID collision
LN -- local site not designated  LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status  IL -- no incoming label
MM -- MTU mismatch               MI -- Mesh-Group ID not available
BK -- Backup connection          ST -- Standby connection
PF -- Profile parse failure      PB -- Profile busy
RS -- remote site standby        SN -- Static Neighbor
LB -- Local site not best-site   RB -- Remote site not best-site
VM -- VLAN ID mismatch

Legend for interface status
Up -- operational
Dn -- down

Instance: vpls
  VPLS-id: 100
  Mesh-group connections: __ves__
    Neighbor                  Type  St     Time last up          # Up trans
    10.255.114.3 (vpls-id 100)  rmt   Up     Apr 11 23:38:38 2013           1
      Remote PE: 10.255.114.3, Negotiated control-word: No
      Incoming label: 262145, Outgoing label: 262145
      Negotiated PW status TLV: No
      Local interface: lsi.1049090, Status: Up, Encapsulation: ETHERNET
        Description: Intf - vpls h-vpls neighbor 10.255.114.3 vpls-id 100

  Mesh-group connections: spokes
    Neighbor                  Type  St     Time last up          # Up trans
    10.255.114.4 (vpls-id 200)  rmt   Up     Apr 11 23:39:25 2013           1
      Remote PE: 10.255.114.4, Negotiated control-word: No
      Incoming label: 262148, Outgoing label: 304224
      Negotiated PW status TLV: Yes
      local PW status code: 0x00000000, Neighbor PW status code: 0x00000000
      Local interface: lsi.1049091, Status: Up, Encapsulation: ETHERNET
        Description: Intf - vpls h-vpls neighbor 10.255.114.4 vpls-id 200
      10.255.114.4 (vpls-id 201)  rmt   Up     Apr 11 23:39:25 2013           1
        Remote PE: 10.255.114.4, Negotiated control-word: No
        Incoming label: 262149, Outgoing label: 304225
        Negotiated PW status TLV: Yes
        local PW status code: 0x00000000, Neighbor PW status code: 0x00000000
        Local interface: lsi.1049096, Status: Up, Encapsulation: ETHERNET
        Description: Intf - vpls h-vpls neighbor 10.255.114.4 vpls-id 201
```

Copyright © 2018, Juniper Networks, Inc.

show vpls connections extensive (Static VPLS Neighbors)

```
user@host> show vpls connections extensive instance red
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down    NP -- interface hardware not present
CM -- control-word mismatch      -> -- only outbound connection is up
CN -- circuit not provisioned    <- -- only inbound connection is up
OR -- out of range              Up -- operational
OL -- no outgoing label          Dn -- down
LD -- local site signaled down   CF -- call admission control failure
RD -- remote site signaled down  SC -- local and remote site ID collision
LN -- local site not designated  LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unn connection status  IL -- no incoming label
MM -- MTU mismatch              MI -- Mesh-Group ID not available
BK -- Backup connection                ST -- Standby connection
PF -- Profile parse failure      PB -- Profile busy
RS -- remote site standby       SN -- Static Neighbor

Legend for interface status
Up -- operational
Dn -- down

Instance: static
  VPLS-id: 1
   Number of local interfaces: 1
   Number of local interfaces up: 1
   ge-0/0/5.0
   lsi.1049344                  Intf - vpls static neighbor 10.255.114.3 vpls-id
 1
   Neighbor                 Type  St     Time last up          # Up trans
   10.255.114.3(vpls-id 1)(SN) rmt Up     Mar  4 08:48:41 2010           1
     Remote PE: 10.255.114.3, Negotiated control-word: No
     Incoming label: 29696, Outgoing label: 29697
     Negotiated PW status TLV: No
     Local interface: lsi.1049344, Status: Up, Encapsulation: ETHERNET
       Description: Intf - vpls static neighbor 10.255.114.3 vpls-id 1
   Connection History:
       Mar  4 08:48:41 2010  status update timer
       Mar  4 08:48:41 2010  PE route changed
       Mar  4 08:48:41 2010  Out lbl Update                    29697
       Mar  4 08:48:41 2010  In lbl Update                     29696
       Mar  4 08:48:41 2010  loc intf up                 lsi.1049344
```

```
user@PE1> show vpls connections extensive (Multihoming with FEC 129)
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down    NP -- interface hardware not present
CM -- control-word mismatch      -> -- only outbound connection is up
CN -- circuit not provisioned    <- -- only inbound connection is up
OR -- out of range               Up -- operational
OL -- no outgoing label          Dn -- down
LD -- local site signaled down   CF -- call admission control failure
```

 Copyright © 2018, Juniper Networks, Inc.

```
RD -- remote site signaled down  SC -- local and remote site ID collision
LN -- local site not designated  LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status  IL -- no incoming label
MM -- MTU mismatch               MI -- Mesh-Group ID not available
BK -- Backup connection          ST -- Standby connection
PF -- Profile parse failure      PB -- Profile busy
RS -- remote site standby        SN -- Static Neighbor
LB -- Local site not best-site   RB -- Remote site not best-site
VM -- VLAN ID mismatch

Legend for interface status
Up -- operational
Dn -- down

Instance: green
  L2vpn-id: 100:100
  Local-id: 192.0.2.2
    Number of local interfaces: 2
    Number of local interfaces up: 2
    ge-0/3/1.0
    ge-0/3/3.0
    lsi.101711873                 Intf - vpls green local-id 192.0.2.2 remote-id
 192.0.2.4 neighbor 192.0.2.4
    Remote-id                 Type  St     Time last up          # Up trans
    192.0.2.4                  rmt   Up     Jan 31 13:49:52 2012           1
      Remote PE: 192.0.2.4, Negotiated control-word: No
      Incoming label: 262146, Outgoing label: 262146
      Local interface: lsi.101711873, Status: Up, Encapsulation: ETHERNET
        Description: Intf - vpls green local-id 192.0.2.2 remote-id 192.0.2.4
neighbor 192.0.2.4
    Connection History:
        Jan 31 13:49:52 2012  status update timer
        Jan 31 13:49:52 2012  PE route changed
        Jan 31 13:49:52 2012  Out lbl Update                    262146
        Jan 31 13:49:52 2012  In lbl Update                     262146
        Jan 31 13:49:52 2012  loc intf up                lsi.101711873
  Multi-home:
    Local-site                Id     Pref   State
    test                      1      100    Up
      Number of interfaces: 1
      Number of interfaces up: 1
        ge-0/3/1.0
      Received multi-homing advertisements:
        Remote-PE         Pref   flag   Description
        192.0.2.4           100    0x0
```

show vpls flood event-queue

Syntax show vpls flood event-queue

Release Information Command introduced in Junos OS Release 8.0.

Description Display the pending events in the VPLS flood queue.

Options This command has no options.

Required Privilege Level view

List of Sample Output show vpls flood event-queue on page 2544

Output Fields Table 243 on page 2544 lists the output fields for the **show vpls flood event-queue** command. Output fields are listed in the approximate order in which they appear.

Table 243: show vpls flood event-queue Output Fields

Field Name	Field Description
Current Pending Event	Provides information on the current event in the VPLS flood event queue.
Name	Name of the event.
Owner Name	Name of the interface associated with the flood event.
Pending Op	Pending operation for the event.
Last Error	Name of the last error encountered.
Number of Retries	Number of attempts made to update the event queue.
Pending Event List	List of the events awaiting processing.
Event Name	Name of the event.
Pending Op	Pending operation for the event.
Event Identifier	Name of the interface associated with the flood event.

Sample Output

show vpls flood event-queue

`user@host>` **show vpls flood event-queue**

Copyright © 2018, Juniper Networks, Inc.

```
Current Pending Event
    Name:          Flood Nexthop
    Owner Name:ge-4/3/0.0
    Pending Op: ADD
    Last Error:ENOMEM
    Number of Retries:3
    Pending Event List:
    Event Name      Pending Op      Event Identifier
    Flood Nexthop   ADD             ge-4/3/0.0
    Flood Route     ADD             ge-4/3/0.0
```

Copyright © 2018, Juniper Networks, Inc.

show vpls flood instance

Syntax

show vpls flood instance
<brief | detail | extensive>
<*instance-name*>
<logical-system *logical-system-name*>

Release Information Command introduced in Junos OS Release 8.0.

Description Display VPLS information related to the flood process.

Options

none—Display VPLS information related to the flood process for all routing instances.

brief | detail | extensive—(Optional) Display the specified level of output.

instance-name—(Optional) Display VPLS information related to the flood process for the specified routing instance.

logical-system *logical-system-name*—(Optional) Display VPLS information related to the flood process for the specified logical system.

Required Privilege Level view

List of Sample Output

show vpls flood instance on page 2547
show vpls flood instance logical-system-name on page 2547
show vpls flood instance detail on page 2547

Output Fields Table 244 on page 2546 lists the output fields for the **show vpls flood instance** command. Output fields are listed in the approximate order in which they appear.

Table 244: show vpls flood instance Output Fields

Field Name	Field Description
Logical system	Name of the logical system.
Name	Name of the VPLS routing instance.
CEs	Number of CE routers connected to the VPLS instance.
VEs	Number of VE routers connected to the VPLS instance.
Flood routes	List of all flood routes associated with the VPLS instance.
Prefix	Prefix for the route.
Type	Type of route.

 Copyright © 2018, Juniper Networks, Inc.

Table 244: show vpls flood instance Output Fields (continued)

Field Name	Field Description
Owner	VPLS routing instance or interface associated with the route.
Nhtype	Next-hop type. For example, **flood** for a flood route.
Nhindex	Next-hop index number for the route.

Sample Output

show vpls flood instance

```
user@host> show vpls flood instance

Logical system: __example_ls1__
Name: green
CEs: 1
VEs: 1
Flood Routes:
  Prefix     Type          Owner                  NhType          NhIndex
  default    ALL_CE_FLOOD  green                  flood           383
  0x47/16    CE_FLOOD      fe-1/2/1.0             flood           388
```

show vpls flood instance logical-system-name

```
user@host:__example_ls1__> show vpls flood instance example_ls1


Logical system: __example_ls1__
Name: green
CEs: 1
VEs: 1
Flood Routes:
  Prefix     Type          Owner                  NhType          NhIndex
  default    ALL_CE_FLOOD  green                  flood           383
  0x47/16    CE_FLOOD      fe-1/2/1.0             flood           388
```

show vpls flood instance detail

```
user@host:__example_ls1__> show vpls flood instance detail


Logical system: __example_ls1__
Name: green
CEs: 1
VEs: 1
Flood Routes:
  Prefix     Type          Owner                  NhType          NhIndex
  default    ALL_CE_FLOOD  green                  flood           383
  0x47/16    CE_FLOOD      fe-1/2/1.0             flood           388
```

Copyright © 2018, Juniper Networks, Inc.

show vpls flood route

Syntax

show vpls flood route
(all-ce-flood instance-name *instance-name* <logical-system-name *logical-system-name*>
  |
ce-flood interface *interface-name*)

Release Information

Command introduced in Junos OS Release 8.0.

Description

Display VPLS route information related to the flood process for either the specified routing instance or the specified interface.

Options

all-ce-flood—Display the flood next-hop route for all customer edge routers for traffic coming from the core of the network.

ce-flood interface ***interface-name***—Display the flood next-hop route for traffic coming from the specified customer edge interface.

instance-name ***instance-name***—Display the flood routes for the specified instance.

logical-system-name ***logical-system-name***—(Optional) Specify the logical system whose flood routes you want to display. You can only specify the default logical system name for VPLS. The default logical system name is **__example_ls1__**\ (the name must be entered in the command with the underscore characters).

Required Privilege Level

view

List of Sample Output

show vpls flood route all-ce-flood on page 2549
show vpls flood route ce-flood on page 2549

Output Fields

Table 245 on page 2548 lists the output for the **show vpls flood route** command. Output fields are listed in the approximate order in which they appear.

Table 245: show vpls flood route Output Fields

Field Name	Field Description
Flood route prefix	Prefix for the flood route.
Flood route type	Type of flood route (either **CE_FLOOD** or **ALL_CE_FLOOD**).
Flood route owner	VPLS routing instance or interface associated with the flood route.
Nexthop type	Next-hop type. For example, **flood** for a flood route.
Nexthop index	Next-hop index number for the route.
Interfaces flooding to	Interfaces to which VPLS routes are being flooded.

Copyright © 2018, Juniper Networks, Inc.

Table 245: show vpls flood route Output Fields (continued)

Field Name	Field Description
Name	Name of the interface.
Type	Type of VPLS router (**CE** or **VE**).
Nh type	Next-hop type.
Index	Index number for the flood route.

Sample Output

show vpls flood route all-ce-flood

`user@host:__example_ls1__>` **show vpls flood route all-ce-flood logical-system-name __example_ls1__instance-name green**

```
  Flood route prefix: default
  Flood route type: ALL_CE_FLOOD
  Flood route owner: green
  Nexthop type: flood
  Nexthop index: 383
    Interfaces Flooding to:
    Name             Type          NhType          Index
    fe-1/2/1.0       CE
```

show vpls flood route ce-flood

`user@host:__example_ls1__>` **show vpls flood route ce-flood interface fe-1/2/1.0**

```
  Flood route prefix: 0x47/16
  Flood route type: CE_FLOOD
  Flood route owner: fe-1/2/1.0
  Nexthop type: flood
  Nexthop index: 388
    Interfaces Flooding to:
    Name             Type          NhType          Index
    lsi.49152        VE             indr            262142
```

Copyright © 2018, Juniper Networks, Inc.

show vpls mac-table

Syntax

show vpls mac-table
<age>
<brief | detail | extensive | summary>
<bridge-domain *bridge-domain-name*>
<instance *instance-name*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>
<*mac-address*>
<vlan-id *vlan-id-number*>

Release Information

Command introduced in Junos OS Release 8.5.
Command introduced in Junos OS Release 15.1.

Description

Display learned virtual private LAN service (VPLS) media access control (MAC) address information.

Options

none—Display all learned VPLS MAC address information.

age— (Optional) Display age of a single mac-address.

brief | detail | extensive | summary—(Optional) Display the specified level of output.

bridge-domain *bridge-domain-name*—(Optional) Display learned VPLS MAC addresses for the specified bridge domain.

instance *instance-name*—(Optional) Display learned VPLS MAC addresses for the specified instance.

interface *interface-name*—(Optional) Display learned VPLS MAC addresses for the specified instance.

logical-system (all | *logical-system-name*)—(Optional) Display learned VPLS MAC addresses for all logical systems or for the specified logical system.

mac-address—(Optional) Display the specified learned VPLS MAC address information..

vlan-id *vlan-id-number*—(Optional) Display learned VPLS MAC addresses for the specified VLAN.

Required Privilege Level

view

List of Sample Output

show vpls mac-table on page 2552
show vpls mac-table (with Layer 2 Services over GRE Interfaces) on page 2552
show vpls mac-table (with VXLAN enabled) on page 2552
show vpls mac-table age (for GE interface) on page 2553
show vpls mac-table age (for AE interface) on page 2553
show vpls mac-table count on page 2553

Copyright © 2018, Juniper Networks, Inc.

show vpls mac-table detail on page 2554
show vpls mac-table extensive on page 2554

Output Fields Table 246 on page 2551 describes the output fields for the **show vpls mac-table** command. Output fields are listed in the approximate order in which they appear.

Table 246: show vpls mac-table Output fields

Field Name	Field Description
Age	Age of a single mac-address.
Routing instance	Name of the routing instance.
Bridging domain	Name of the bridging domain.
MAC address	MAC address or addresses learned on a logical interface.
MAC flags	Status of MAC address learning properties for each interface: • **S**—Static MAC address configured. • **D**—Dynamic MAC address learned. • **SE**—MAC accounting is enabled. • **NM**—Nonconfigured MAC.
Logical interface	Name of the logical interface.
MAC count	Number of MAC addresses learned on a specific routing instance or interface.
Learning interface	Logical interface or logical Label Switched Interface (LSI) the address is learned on.
Base learning interface	Base learning interface of the MAC address. This field is introduced in Junos OS Release 14.2.
Learn VLAN ID/VLAN	VLAN ID of the routing instance or bridge domain in which the MAC address was learned.
VXLAN ID/VXLAN	VXLAN Network Identifier (VNI)
Layer 2 flags	Debugging flags signifying that the MAC address is present in various lists.
Epoch	Spanning Tree Protocol epoch number identifying when the MAC address was learned. Used for debugging.
Sequence number	Sequence number assigned to this MAC address. Used for debugging.
Learning mask	Mask of Packet Forwarding Engines where this MAC address was learned. Used for debugging.
IPC generation	Creation time of the logical interface when this MAC address was learned. Used for debugging.

Copyright © 2018, Juniper Networks, Inc.

Sample Output

show vpls mac-table

```
user@host> show vpls mac-table
MAC flags (S -static MAC, D -dynamic MAC,
           SE -Statistics enabled, NM -Non configured MAC)

Routing instance : vpls_ldp1
VLAN : 223
   MAC                  MAC      Logical
   address              flags    interface
   00:00:5e:00:53:5d    D        ge-0/2/5.400

MAC flags (S -static MAC, D -dynamic MAC,
           SE -Statistics enabled, NM -Non configured MAC)

Routing instance : vpls_red
VLAN : 401
   MAC                  MAC      Logical
   address              flags    interface
   00:00:5e:00:53:12    D        lsi.1051138
   00:00:5e:00:53:f0    D        lsi.1051138
```

show vpls mac-table (with Layer 2 Services over GRE Interfaces)

```
user@host> show vpls mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned
           SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : vpls_4site:1000
 Bridging domain : __vpls_4site:1000__,   MAC                   MAC      Logical
   address              flags    interface
   00:00:5e:00:53:f4    D,SE     ge-4/2/0.1000
   00:00:5e:00:53:33    D,SE     lsi.1052004
   00:00:5e:00:53:32    D,SE     lsi.1048840
   00:00:5e:00:53:14    D,SE     lsi.1052005
   00:00:5e:00:53:f7    D,SE     gr-1/2/10.10
```

show vpls mac-table (with VXLAN enabled)

```
user@host> show vpls mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned
           SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : vpls_4site:1000
 Bridging domain : __vpls_4site:1000__, VLAN : 4094,4093
  VXLAN: Id : 300, Multicast group: 233.252.0.1
   MAC                  MAC      Logical
   address              flags    interface
   00:00:5e:00:53:f4    D,SE     ge-4/2/0.1000
   00:00:5e:00:53:33    D,SE     lsi.1052004
   00:00:5e:00:53:32    D,SE     lsi.1048840
   00:00:5e:00:53:14    D,SE     lsi.1052005
   00:00:5e:00:53:f7    D,SE     vtep.1052010
   00:00:5e:00:53:3f    D,SE     vtep.1052011
```

 Copyright © 2018, Juniper Networks, Inc.

show vpls mac-table age (for GE interface)

```
user@host> show vpls mac-table age 00:00:5e:00:53:1a instance vpls_instance_1
MAC Entry Age information
Current Age: 4 seconds
```

show vpls mac-table age (for AE interface)

```
user@host> show vpls mac-table age 000:00:5e:00:53:1a instance vpls_instance_1
MAC Entry Age information
Current Age on FPC1: 102 seconds
Current Age on FPC2: 94 seconds
```

show vpls mac-table count

```
user@host> show vpls mac-table count
0 MAC address learned in routing instance __example_private1__

  MAC address count per interface within routing instance:
    Logical interface        MAC count
    lc-0/0/0.32769                   0
    lc-0/1/0.32769                   0
    lc-0/2/0.32769                   0
    lc-2/0/0.32769                   0
    lc-0/3/0.32769                   0
    lc-2/1/0.32769                   0
    lc-9/0/0.32769                   0
    lc-11/0/0.32769                  0
    lc-2/2/0.32769                   0
    lc-9/1/0.32769                   0
    lc-11/1/0.32769                  0
    lc-2/3/0.32769                   0
    lc-9/2/0.32769                   0
    lc-11/2/0.32769                  0
    lc-11/3/0.32769                  0
    lc-9/3/0.32769                   0

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID            MAC count
                0                    0

1 MAC address learned in routing instance vpls_ldp1

  MAC address count per interface within routing instance:
    Logical interface        MAC count
    lsi.1051137                      0
    ge-0/2/5.400                     1

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID            MAC count
                0                    1

1 MAC address learned in routing instance vpls_red

  MAC address count per interface within routing instance:
    Logical interface        MAC count
    ge-0/2/5.300                     1

  MAC address count per learn VLAN within routing instance:
```

```
    Learn VLAN ID             MAC count
                0                     1
```

show vpls mac-table detail

```
user@host> show vpls mac-table detail
MAC address: 00:00:5e:00:53:5d
  Routing instance: vpls_ldp1
   Learning interface: ge-0/2/5.400
   Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
   Epoch: 0                            Sequence number: 1
   Learning mask: 0x1                  IPC generation: 0

MAC address: 00:00:5e:00:53:5d
  Routing instance: vpls_red
   Learning interface: ge-0/2/5.300
   Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
   Epoch: 0                            Sequence number: 1
   Learning mask: 0x1                  IPC generation: 0
```

show vpls mac-table extensive

```
user@host> show vpls mac-table extensive

MAC address: 00:00:5e:00:53:00
  Routing instance: vpls_1
   Bridging domain: __vpls_1__, VLAN : NA
   Learning interface: lsi.1049165
   Base learning interface: lsi.1049165
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                            Sequence number: 1
   Learning mask: 0x00000001

MAC address: 00:00:5e:00:53:01
  Routing instance: vpls_1
   Bridging domain: __vpls_1__, VLAN : NA
   Learning interface: lsi.1049165
   Base learning interface: lsi.1049165
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                            Sequence number: 1
   Learning mask: 0x00000001

MAC address: 00:00:5e:00:53:02
  Routing instance: vpls_1
   Bridging domain: __vpls_1__, VLAN : NA
   Learning interface: lsi.1049165
   Base learning interface: lsi.1049165
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                            Sequence number: 1
   Learning mask: 0x00000001

MAC address: 00:00:5e:00:53:03
  Routing instance: vpls_1
   Bridging domain: __vpls_1__, VLAN : NA
   Learning interface: lsi.1049165
   Base learning interface: lsi.1049165
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                            Sequence number: 1
   Learning mask: 0x00000001
```

 Copyright © 2018, Juniper Networks, Inc.

show vpls statistics

Syntax

show vpls statistics
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced before Junos OS Release 7.4.

Description

(T Series and M Series routers, except for the M160 router) Display virtual private LAN service (VPLS) statistics.

Options

none—Display VPLS statistics for all routing instances.

instance *instance-name*—(Optional) Display VPLS statistics for a specific VPLS routing instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show vpls statistics on page 2556
show vpls statistics instance on page 2557

Output Fields

Table 247 on page 2555 lists the output fields for the **show vpls statistics** command. Output fields are listed in the approximate order in which they appear.

Table 247: show vpls statistics Output Fields

Field Name	Field Description
Instance	Name of the VPLS instance.
Local interface	Name of the local VPLS virtual loopback tunnel interface, **vt-*fpc/pic/port.nnnnn***, where ***nnnnn*** is a dynamically generated virtual port used to transport and receive packets from other provider edge (PE) routers in the VPLS domain.
Index	Number associated with the next hop.
Remote provider edge router	Address of the remote provider edge router.
Multicast packets	Number of multicast packets received.
Multicast bytes	Number of multicast bytes received.
Flood packets	Number of VPLS flood packets received.

Copyright © 2018, Juniper Networks, Inc.

Table 247: show vpls statistics Output Fields (continued)

Field Name	Field Description
Flood bytes	Number of VPLS flood bytes received.
Current MAC count	Number of MAC addresses learned by the interface and the configured maximum limit on the number of MAC addresses that can be learned.

Sample Output

show vpls statistics

user@host> **show vpls statistics**

```
VPLS statistics:

Instance: green

   Local interface: fe-2/2/1.0, Index: 69
     Multicast packets:                     1
     Multicast bytes  :                    60
     Flooded packets  :                    18
     Flooded bytes    :                  2556
     Current MAC count:                     1

   Local interface: lt-0/3/0.2, Index: 72
     Multicast packets:                     3
     Multicast bytes  :                   153
     Flooded packets  :                     1
     Flooded bytes    :                    51
     Current MAC count:                     1

   Local interface: lsi.32769, Index: 75
     Current MAC count:                     0

   Local interface: lsi.32771, Index: 77
   Remote PE: 10.255.14.222
     Current MAC count:                     2

Instance: red

   Local interface: vt-0/3/0.32768, Index: 74
     Multicast packets:                     0
     Multicast bytes  :                     0
     Flooded packets  :                     0
     Flooded bytes    :                     0
     Current MAC count:                     0

   Local interface: vt-0/3/0.32770, Index: 76
     Multicast packets:                     0
     Multicast bytes  :                     0
     Flooded packets  :                     0
     Flooded bytes    :                     0
     Current MAC count:                     0
```


Copyright © 2018, Juniper Networks, Inc.

show vpls statistics instance

user@host> **show vpls statistics instance red**

```
Layer-2 VPN Statistics:
Instance: red

   Local interface: vt-3/2/0.32768, Index: 73
   Remote provider edge router: 10.255.17.35
     Multicast packets:                     0
     Multicast bytes  :                     0
     Flood packets    :                     0
     Flood bytes      :                     0
     Current MAC count:                     1 (Limit 20)
```

Copyright © 2018, Juniper Networks, Inc.