Junos® OS

Triple Play for Subscriber Services Feature Guide

Release
13.2

Published: 2013-07-31

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

Junos® OS Triple Play for Subscriber Services Feature Guide
13.2

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

List of Figures

List of Tables

About the Documentation

Documentation and Release Notes

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

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

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

Supported Platforms

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

- MX Series

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

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

Documentation Conventions

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

Table 1: Notice Icons

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

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

Table 2: Text and Syntax Conventions

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

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

Convention	Description	Examples
Italic text like this	• Introduces or emphasizes important new terms. • Identifies book names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS System Basics Configuration Guide* • RFC 1997, *BGP Communities Attribute*
Italic text like this	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# **set system domain-name** *domain-name*
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	• To configure a stub area, include the **stub** statement at the **[edit protocols ospf area area-id]** hierarchy level. • The console port is labeled **CONSOLE**.
< > (angle brackets)	Enclose optional keywords or variables.	**stub <default-metric** *metric***>;**
| (pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	**broadcast | multicast** (*string1* | *string2* | *string3*)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	**rsvp { # Required for dynamic MPLS only**
[ ] (square brackets)	Enclose a variable for which you can substitute one or more values.	**community name members [** *community-ids* **]**

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

; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
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 send your comments to **techpubs-comments@juniper.net**, or fill out the documentation feedback form at https://www.juniper.net/cgi-bin/docbugreport/ . If you are using e-mail, be sure to include the following information with your comments:

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

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

- Find CSC offerings: http://www.juniper.net/customers/support/
- Search for known bugs: http://www2.juniper.net/kb/
- Find product documentation: http://www.juniper.net/techpubs/
- Find solutions and answer questions using our Knowledge Base: http://kb.juniper.net/
- Download the latest versions of software and review release notes: http://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: https://www.juniper.net/alerts/

- Join and participate in the Juniper Networks Community Forum: http://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: http://www.juniper.net/cm/

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

Opening a Case with JTAC

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

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

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

PART 1

Overview

CHAPTER 1

Broadband Subscriber Management Basics Overview

Broadband Subscriber Management Overview

Broadband Subscriber Management is a method of dynamically provisioning and managing subscriber access in a multiplay or triple play network environment. This method uses AAA configuration in conjunction with dynamic profiles to provide dynamic, per-subscriber authentication, addressing, access, and configuration for a host of broadband services including Internet access, gaming, IPTV, Video on Demand (VoD), and subscriber wholesaling.

NOTE: **The Junos OS broadband subscriber management solution currently supports Dynamic Host Configuration Protocols (DHCP)-based and Point-to-Point Protocol /Point-to-Point Protocol over Ethernet (PPP/PPPoE)-based configuration and RADIUS authentication and authorization.**

This guide focuses on the general components necessary for configuring a Juniper Networks MX Series 3D Universal Edge Router to dynamically provision and manage subscribers. However, you can also use a Juniper Networks EX Series Ethernet Switch in a subscriber network.

Managing subscribers in a DHCP-based or PPP/PPPoE-based residential broadband network using an MX Series router requires the following:

- Planning and configuring a virtual LAN (VLAN) architecture for the access network.
- Configuring an authentication, authorization, and accounting (AAA) framework for subscriber authentication and authorization through external servers (for example, RADIUS) as well as accounting and dynamic-request change of authorization (CoA) and disconnect operations through external servers, and address assignment through a combination of local address-assignment pools and RADIUS.
- Configuring DHCP local server or DHCP relay for subscriber address assignment for DHCP-based networks.
- Configuring address assignment pools for PPPoE-based networks.
- Configuring dynamic profiles to include dynamic IGMP, firewall filter, and class of service (CoS) configuration for subscriber access.
- Configuring multicast access to the core network.

To better understand the subscriber access network, this guide also provides general information about some hardware not from Juniper Networks and suggests methods for choosing different network configuration options. You can configure a subscriber network in many different ways. This guide does not cover all configuration scenarios. It is intended as a starting point for understanding subscriber management and how you can use Juniper Networks hardware and software to plan and build your own subscriber management solution.

Related Documentation

- Broadband Subscriber Management Platform Support on page 4
- Broadband Subscriber Management Network Topology Overview on page 5
- Broadband Subscriber Management Solutions Terms and Acronyms on page 5
- Supporting Documentation for Broadband Subscriber Management on page 7
- Triple Play and Multiplay Overview on page 8
- Broadband History on page 9

Broadband Subscriber Management Platform Support

Juniper Networks currently supports DHCP and PPP/PPPoE broadband subscriber management solutions on MX Series routers and PPP/PPPoE broadband subscriber management solutions on M120 and M320 routers.

NOTE: **This guide describes configuration on MX Series routers.**

Related Documentation

- Broadband Subscriber Management Overview on page 3
- Broadband Subscriber Management Edge Router Overview on page 15

Broadband Subscriber Management Network Topology Overview

Figure 1 on page 5 illustrates how network elements can make up a residential broadband access network.

Figure 1: Subscriber Management Residential Broadband Network Example

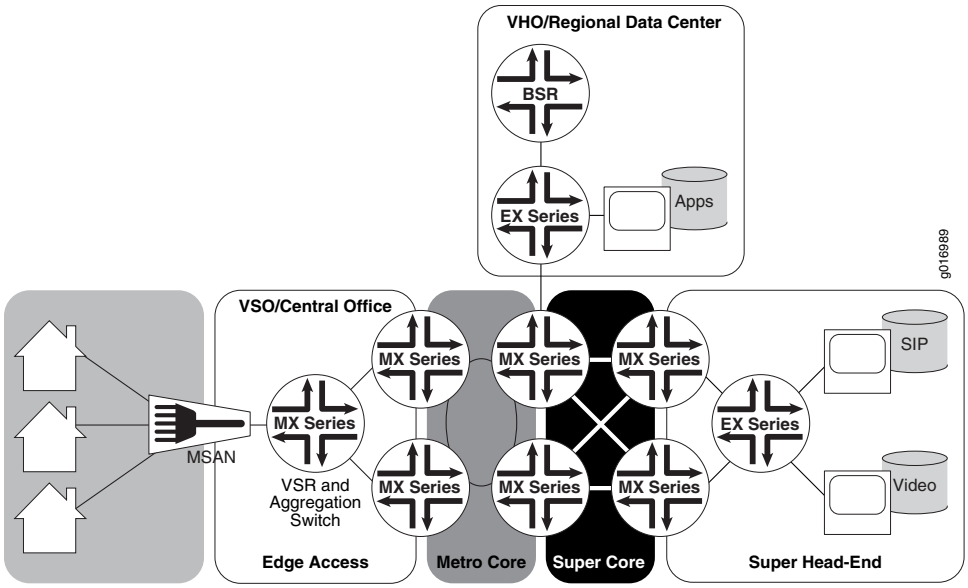

Related Documentation

- Broadband Subscriber Management Overview on page 3

Broadband Subscriber Management Solutions Terms and Acronyms

- **AAA (authentication, authorization, and accounting)**—An IP-based networking system that controls user access to computer resources and manages the activity of users over a network.
- **ASM (Any Source Multicast)**—A method of allowing a multicast receiver to listen to all traffic sent to a multicast group, regardless of its source.
- **BSR (broadband services router)**—A router used for subscriber management and edge routing.
- **CoA (change of authorization)**—RADIUS messages that contain information for dynamically changing session authorizations.
- **CoS (class of service)**—A method of managing network traffic by grouping similar types of traffic together and treating each traffic type as a “class” with a defined service priority.
- **DHCP (Dynamic Host Configuration Protocol)**—A mechanism through which hosts using TCP/IP can obtain protocol configuration parameters automatically from a DHCP

server on the network; allocates IP addresses dynamically so that they can be reused when no longer needed.

- **IGMP (Internet Group Membership Protocol)**—A host-to-router signaling protocol for IPv4 used to support IP multicasting.
- **IS-IS (Intermediate System-to-Intermediate System)**—A link-state interior gateway routing protocol (IGRP) for IP networks that uses the shortest-path-first (SPF) algorithm to determine routes.
- **LSP (label-switched path)**—The path traversed by a packet that is routed by MPLS. Some LSPs act as tunnels. LSPs are unidirectional, carrying traffic only in the downstream direction from an ingress node to an egress node.
- **MPLS (Multiprotocol Label Switching)**—A mechanism for engineering network traffic patterns that functions by assigning to network packets short labels that describe how to forward the packets through the network.
- **MSAN (multiservice access node)**—A group of commonly used aggregation devices including digital subscriber line access multiplexers (DSLAMs) used in xDSL networks, optical line termination (OLT) for PON/FTTx networks, and Ethernet switches for Active Ethernet connections.
- **Multiplay**—A networking paradigm that enables the ability to add new and robust networking services that individual subscribers can access.
- **OIF (outgoing interface)**—An interface used by multicast functions within a router to determine which egress ports to use for fowarding multicast groups.
- **OSPF (Open Shortest Path First)**—A link-state interior gateway protocol (IGP) that makes routing decisions based on the shortest-path-first (SPF) algorithm (also referred to as the Dijkstra algorithm).
- **PIM (Protocol Independent Multicast)**—A multicast routing protocol used for delivering multicast messages in a routed environment.
- **PPP (Point-to-Point Protocol)**—A link-layer protocol that provides multiprotocol encapsulation. PPP is used for link-layer and network-layer configuration. Provides a standard method for transporting multiprotocol datagrams over point-to-point links.
- **PPPoE (Point-to-Point Protocol over Ethernet)**—A network protocol that encapsulates PPP frames in Ethernet frames and connects multiple hosts over a simple bridging access device to a remote access concentrator.
- **RADIUS (Remote Authentication Dial-In User Service)**—A networking protocol that provides centralized access, authorization, and accounting management for subscribers to connect and use a network service.
- **Residential gateway**—A firewall, Network Address Translation (NAT) router, or other routing device used as a customer premises equipment (CPE) terminator in the home, office, or local point of presence (POP).
- **SSM (single-source multicast)**—A routing method that allows a multicast receiver to detect only a specifically identified sender within a multicast group.
- **set-top box**—The end host or device used to receive IPTV video streams.

- **Triple play**—A networking paradigm that dedicates bandwidth to data, voice, and video service.
- **VOD (video on demand)**—A unicast streaming video offering by service providers that enables the reception of an isolated video session per user with rewind, pause, and similar VCR-like capabilities.
- **VSR (video services router)**—A router used in a video services network to route video streams between an access network and a metro or core network. The video services router is any M Series Multiservice Edge Router or MX Series router that supports the video routing package provided with Junos OS Release 8.3 or later.

Related Documentation

- Broadband Subscriber Management Overview on page 3

Supporting Documentation for Broadband Subscriber Management

The *Junos OS Subscriber Management and Services Library* relies heavily on existing configuration documentation. In particular, this guide references configuration material presented in the *Junos OS Subscriber Management and Services Library*. We recommend you become familiar with the configuration options presented for subscriber access before reading this guide.

Several guides in the Junos OS documentation set provide detailed configuration information that is not fully covered in this guide. This guide might reference other Junos OS configuration and solutions documents that can provide more detail about a specific feature or configuration option.

For more detailed configuration information, see the following Junos OS documents:

- *Junos OS Subscriber Management and Services Library*
- *Junos OS Layer 2 Switching and Bridging Library for Routing Devices*
- *Multicast Protocols Feature Guide for Routing Devices*
- *Junos OS Network Interfaces Library for Routing Devices*
- *Routing Policy Feature Guide for Routing Devices*

In addition to related Junos OS documentation, you can obtain useful information from the JunosE Software documentation. Many features described in the *JunosE Broadband Access Configuration Guide* are similar to those described in both this guide and the *Junos OS Subscriber Management and Services Library*.

Related Documentation

- Broadband Subscriber Management Overview on page 3

Triple Play and Multiplay Overview

This document defines triple play and multiplay networks as different entities:

- A *triple play* network dedicates bandwidth to each possible service—data, voice, and video. This method works well when a limited number of services are deployed and sufficient bandwidth is available.
- A *multiplay* network refers to the ability to add new and robust networking services that each subscriber can access. This method requires the integration of dynamic bandwidth management and the ability to manage subscribers dynamically though the use of features such as hierarchical quality of service (QoS) and a AAA service framework that provides authentication, accounting, dynamic change of authorization (CoA), and dynamic address assignment.

Table 3 on page 8 provides some comparison between a triple play and multiplay network and the level of flexibility associated with certain networking options.

Table 3: Triple Play and Multiplay Comparison

Flexibility	Triple Play	Multiplay
Bandwidth Management	Fixed bandwidth allocation for each service.	One bandwidth pool for each subscriber is shared by all services.
Adding New Services	Requires *deallocating* bandwidth from one service and allocating that bandwith to the new service.	The existence of one shared bandwidth pool eliminates the need to reallocate bandwidth to new services.
Subscriber Flexibility	Limited subscriber flexibility because a fixed bandwidth is allocated to each service or application.	Subscribers can use their share of bandwidth for whatever applications they want to run.
Client Device Types	Client devices (PCs or set-top boxes) are dedicated to specific services and often assigned to specific ports on customer premise equipment.	Client devices are not assigned to any specific ports. This flexibility enables the ability to use client devices for various services (for example, adding software to a PC to enable television broadcasts) and allows different client devices (PCs, Voice-over-IP phones, and set-top boxes) to reside on a single LAN.

With software and hardware now available to enable client devices to access and use the network in a variety of ways, bandwidth demands increasing, and new networking business models emerging, dynamic support of new applications is required to ensure subscriber satisfaction. A dynamic multiplay network configuration can provide the flexibility to meet these demands.

Related Documentation

- Broadband Subscriber Management Overview on page 3

CHAPTER 2

Residential Broadband Technology Overview

Broadband History

Residential broadband services developed using a mainly ATM-based infrastructure and early Internet access required that each subscriber access the network using a dial-up modem to connect from a PC to a Remote Access Server (RAS), or bank of servers, which was connected directly to the Internet. Point-to-Point Protocol (PPP), originally defined by the IETF in RFC 1661, was already in use on leased lines. It was well suited for use on the existing ATM infrastructure and enabled operators to better manage subscriber connections by providing authentication and accounting, along with a level of protocol flexibility due to it being connection-oriented and enabling service providers to customize it to their needs. The use of the PPP model, however, required special software (including the PPP protocol stack) be installed on each PC to communicate within the PPP network. After establishing a connection to the Internet, the subscriber logged in using a PPP user identifier provided by the service provider.

This *always on* model quickly evolved in several ways. Dedicated *broadband* access such as DSL replaced dial-up service, replacing the dial-up modem with a DSL modem. Dial-up remote access servers were replaced by the Broadband Remote Access Server (B-RAS) and residential gateways were introduced to allow multiple PCs from one site to connect to the broadband network. Residential gateways have since evolved to provide a wide range of functions including firewall and wireless (802.1b/g/n wi-fi) connectivity. The residential gateway also became the termination point for the PPP connection, eliminating the need for the installation of special PC software.

These new broadband networks were built based on the following two key assumptions:

- Only a small percentage of subscribers were expected to be using network bandwidth at any given time and, even if many subscribers logged in to the network concurrently, few subscribers were likely to enter data at the exact same time.

- Traffic was TCP-based and not real-time. If a packet was lost due to network congestion, TCP detected the loss and retransmitted the packets.

Based on these assumptions, operators over-subscribed the network, enabling more subscribers than a limited amount of bandwidth can support if all subscribers were to access the network simultaneously. For example, if 50 subscribers were to sign up for service that required bandwidth of 1 Mbps for each subscriber, the network did not necessarily need to support a full 50 Mbps of throughput. Instead, operators designed the network to support much lower traffic volumes, expecting maximum traffic flow for all subscribers to occur rarely, if ever. For example, a 50:1 over-subscription needed to support only 1 Mbps of bandwidth. Bandwidth requirements have changed significantly over the years and this method of access is becoming more difficult to maintain.

The basic broadband architecture was initially defined by DSL Forum TR-025 (November 1999). This specification assumed only one service was provided to subscribers—Internet Access (or *data*). DSL Forum TR-059 (September 2003) introduced quality of service (QoS) to allow broadband networks to deliver voice over IP (VoIP) in addition to data. Because VoIP is a small percentage of overall network traffic, its introduction has not significantly altered the broadband delivery landscape. It is also worth noting that these original standards specified ATM as the Layer 2 protocol on the broadband network.

Related Documentation

- *PPP in Broadband Networks*
- DHCP in Broadband Networks on page 10
- Broadband Service Delivery Options on page 10
- Broadband Delivery and FTTx on page 12

DHCP in Broadband Networks

Dynamic Host Configuration Protocol (DHCP) is an alternative to PPP for assigning IP addresses and provisioning services in broadband networks. Using DHCP helps to simplify network configuration by decreasing (and in some cases eliminating) the need for manually configuring static IP addresses on network devices. For example, DHCP enables PCs and other devices within a subscriber residence to obtain IP addresses to access the Internet. Due to its general simplicity and scalability, along with the increased usage of Ethernet in access networks, DHCP deployments in broadband networks have increased.

Related Documentation

- Broadband Service Delivery Options on page 10

Broadband Service Delivery Options

Four primary delivery options exist today for delivering broadband network service. These options include the following:

- Digital Subscriber Line
- Active Ethernet

- Passive Optical Networking
- Hybrid Fiber Coaxial

The following sections briefly describe each delivery option.

Digital Subscriber Line

Digital subscriber line (DSL) is the most widely deployed broadband technology worldwide. This delivery option uses existing telephone lines to send broadband information on a different frequency than is used for the existing voice service. Many generations of DSL are used for residential service, including Very High Speed Digital Subscriber Line 2 (VDSL2) and versions of Asymmetric Digital Subscriber Line (ADSL, ADSL2, and ADSL2+). These variations of DSL primarily offer asymmetric residential broadband service where different upstream and downstream speeds are implemented. (VDSL2 also supports symmetric operation.) Other DSL variations, like High bit rate Digital Subscriber Line (HDSL) and Symmetric Digital Subscriber Line (SDSL), provide symmetric speeds and are typically used in business applications.

The head-end to a DSL system is the Digital Subscriber Line Access Multiplexer (DSLAM). The demarcation device at the customer premise is a DSL modem. DSL service models are defined by the Broadband Forum (formerly called the DSL Forum).

Active Ethernet

Active Ethernet uses traditional Ethernet technology to deliver broadband service across a fiber-optic network. Active Ethernet does not provide a separate channel for existing voice service, so VoIP (or TDM-to-VoIP) equipment is required. In addition, sending full-speed (10 or 100 Mbps) Ethernet requires significant power, necessitating distribution to Ethernet switches and optical repeaters located in cabinets outside of the central office. Due to these restrictions, early Active Ethernet deployments typically appear in densely populated areas.

Passive Optical Networking

Passive Optical Networking (PON), like Active Ethernet, uses fiber-optic cable to deliver services to the premises. This delivery option provides higher speeds than DSL but lower speeds than Active Ethernet. Though PON provides higher speed to each subscriber, it requires a higher investment in cable and connectivity.

A key advantage of PON is that it does not require any powered equipment outside of the central office. Each fiber leaving the central office is split using a non-powered optical splitter. The split fiber then follows a point-to-point connection to each subscriber.

PON technologies fall into three general categories:

- ATM PON (APON), Broadband PON (BPON), and Gigabit-capable PON (GPON)—PON standards that use the following different delivery options:
  - APON—The first passive optical network standard is primarily used for business applications.

  - BPON—Based on APON, BPON adds wave division multiplexing (WDM), dynamic and higher upstream bandwidth allocation, and a standard management interface to enable mixed-vendor networks.
  - GPON—The most recent PON adaptation, GPON is based on BPON but supports higher rates, enhanced security, and a choice of which Layer 2 protocol to use (ATM, Generic Equipment Model [GEM], or Ethernet).

- Ethernet PON (EPON)—Provides capabilities similar to GPON, BPON, and APON, but uses Ethernet standards. These standards are defined by the IEEE. Gigabit Ethernet PON (GEPON) is the highest speed version.
- Wave Division Multiplexing PON (WDM-PON)—A nonstandard PON which, as the name implies, provides a separate wavelength to each subscriber.

The head-end to a PON system is an Optical Line Terminator (OLT). The demarcation device at the customer premises is an Optical Network Terminator (ONT). The ONT provides subscriber-side ports for connecting Ethernet (RJ-45), telephone wires (RJ-11) or coaxial cable (F-connector).

Hybrid Fiber Coaxial

Multi-System Operators (MSOs; also known as *cable TV operators*) offer broadband service through their hybrid fiber-coaxial (HFC) network. The HFC network combines optical fiber and coaxial cable to deliver service directly to the customer. Services leave the central office (CO) using a fiber-optic cable. The service is then converted outside of the CO to a coaxial cable *tree* using a series of optical nodes and, where necessary, through a trunk radio frequency (RF) amplifier. The coaxial cables then connect to multiple subscribers. The demarcation device is a cable modem or set-top box, which talks to a Cable Modem Termination System (CMTS) at the MSO *head-end* or master facility that receives television signals for processing and distribution. Broadband traffic is carried using the Data Over Cable Service Interface Specification (DOCSIS) standard defined by CableLabs and many contributing companies.

Related Documentation

- Broadband Delivery and FTTx on page 12

Broadband Delivery and FTTx

Many implementations use existing copper cabling to deliver signal to the premises, but fiber-optic cable connectivity is making its way closer to the subscriber. Most networks use a combination of both copper and fiber-optic cabling. The term *fiber to the x* (FTTx) describes how far into the network fiber-optic cabling runs before a switch to copper cabling takes place. Both PON and Active Ethernet can use fiber-optic portion of the network, while xDSL is typically used on the copper portion. This means that a single fiber-optic strand may support multiple copper-based subscribers.

Increasing the use of fiber in the network increases cost but it also increases network access speed to each subscriber.

The following terms are used to describe the termination point of fiber-optic cable in a network:

- Fiber to the Premises (FTTP), Fiber to the Home (FTTH), Fiber to the Business (FTTB)—Fiber extends all the way to the subscriber. PON is most common for residential access, although Active Ethernet can be efficiently used in dense areas such as apartment complexes. Active Ethernet is more common for delivering services to businesses.
- Fiber to the Curb (FTTC)—Fiber extends most of the way (typically, 500 feet/150 meters or less) to the subscriber. Existing copper is used for the remaining distance to the subscriber.
- Fiber to the Node/Neighborhood (FTTN)—Fiber extends to within a few thousand feet of the subscriber and converted to xDSL for the remaining distance to the subscriber.
- Fiber to the Exchange (FTTE)—A typical central office-based xDSL implementation in which fiber is used to deliver traffic to the central office and xDSL is used on the existing local loop.

Related Documentation

- Broadband Service Delivery Options on page 10

CHAPTER 3

Broadband Subscriber Management Solution Hardware Overview

Broadband Subscriber Management Edge Router Overview

The edge router is the demarcation point between the residential broadband access network and the core network. The Juniper Networks MX Series router (along with the Juniper Networks EX Series Ethernet Switch) can play multiple roles as an edge router. The most common include the following:

- **Broadband services router (BSR)**—This router supports high speed Internet access along with several other subscriber-based services including VoIP, IPTV, and gaming.
- **Video services router (VSR)**—The video services router capabilities are a subset of those provided by a broadband services router. In general, using the MX Series router as a video services router provides bi-directional traffic destined for the set-top box (STB). This traffic includes IPTV and video on demand (VoD) streams as well as associated control traffic such as IGMP and electronic program guide (EPG) updates.

You can also use the MX Series router in certain Layer 2 solutions. For information about configuring the MX Series router in Layer 2 scenarios, see the *Junos OS Layer 2 Switching and Bridging Library for Routing Devices* or the *Ethernet Networking Feature Guide for MX Series Routers*.

Broadband Services Router Overview

A broadband services router is an edge router that traditionally supports primarily Internet-bound traffic. This router replaces and provides a superset of the functionality provided by a Broadband Remote Access Server (B-RAS). The broadband services router functions can be broken into two key areas—high speed Internet access and IPTV support.

High-Speed Internet Access Support

The broadband services router communicates with the RADIUS server to enforce which services each subscriber can access. For example, one subscriber might have signed up

for a smaller Internet access service of 1 Mbps where another subscriber might have signed up for a higher, 10 Mbps service. The broadband services router manages the traffic to each subscriber, ensuring that each subscriber obtains the level of access service they have purchased, while also ensuring that any VoIP traffic receives priority. The broadband services router also makes traffic forwarding decisions based on aggregate bandwidth detected on any adjacent multiservice access node (MSAN).

IPTV Support

The broadband services router supports IPTV traffic including support for IGMP multicast group start and stop requests from downstream MSANs. The broadband services router manages the bandwidth allocations associated with high-bandwidth IPTV as well as video on demand (VoD) traffic to ensure high quality service delivery.

Video Services Router

When configuring a multiedge network, you can use the MX Series router as a video services router (VSR) to support only video traffic without supporting the high-speed Internet access (HSIA) capabilities.

NOTE: **We recommend a single-edge network model but the MX Series router allows for flexibility when defining a multiplay network topology.**

Some advantages of using a separate video services router for video traffic include the following:

- Provides the ability to add IPTV service without the need to modify an existing edge router that is performing other functions.
- Reduces network bandwidth by moving the video edge further out to the network edge while still allowing for centralized broadband services router operation.
- Typically requires less capital investment because the video services router does not need to provide per-subscriber management.

Services Router Placement

Depending on the type of network you are creating—single edge or multiedge—you can place a broadband services router or video services router in various locations.

Single-Edge Placement

In a single-edge network, you use only broadband services routers because the single device must perform all of the necessary edge functions—providing subscriber management for high-speed Internet access and IPTV services. You can use the two following topology models when placing the broadband services router:

- **Centralized single edge**—The edge router is centrally located and placed at one location to cover a particular region. A secondary router is sometimes placed in this location to act as a backup. Downstream MSANs are connected to the broadband services router using a ring or mesh topology.

- **Distributed single edge**—The edge router is placed further out into the network, typically in the central office (CO) closest to the subscribers that it services. Downstream MSANs are typically connected directly to the broadband services router (in a true, single edge topology) or through an Ethernet aggregation switch.

In general, the addition of IPTV service favors a more distributed model because it pushes the need for subscriber management farther out into the network.

Multiedge Placement

In a multiedge network, you use both broadband services routers and video services routers. The broadband services router controls any high-speed Internet traffic and the video services router controls video traffic. You can use the two following topology models when placing service routers in a multiedge network topology:

- **Co-located multiedge**—The broadband services router and video services router are housed in the same location and an Ethernet switch directs traffic in the CO to the appropriate edge router.

  NOTE: A single MX Series router can serve as both Ethernet switch and video services router. For information about configuring the MX Series router in Layer 2 scenarios, see the *Junos OS Layer 2 Switching and Bridging Library for Routing Devices* or the *Ethernet Networking Feature Guide for MX Series Routers*.

- **Split multiedge**—The video services router and broadband services router reside in different locations. In this model, the broadband services router is typically located more centrally and video services routers are distributed.

Related Documentation

- Multiservice Access Node Overview on page 17
- Ethernet MSAN Aggregation Options on page 19
- Broadband Subscriber Management Platform Support on page 4

Multiservice Access Node Overview

A *multiservice access node* is a broader term that refers to a group of commonly used aggregation devices. These devices include digital subscriber line access multiplexers (DSLAMs) used in xDSL networks, optical line termination (OLT) for PON/FTTx networks, and Ethernet switches for Active Ethernet connections. Modern MSANs often support all of these connections, as well as providing connections for additional circuits such as plain old telephone service (referred to as POTS) or Digital Signal 1 (DS1 or T1).

The defining function of a multiservice access node is to aggregate traffic from multiple subscribers. At the physical level, the MSAN also converts traffic from the *last mile technology* (for example, ADSL) to Ethernet for delivery to subscribers.

You can broadly categorize MSANs into three types based on how they forward traffic in the network:

- **Layer–2 MSAN**—This type of MSAN is essentially a Layer 2 switch (though typically not a fully functioning switch) with some relevant enhancements. These MSANs use Ethernet (or ATM) switching to forward traffic. The MSAN forwards all subscriber traffic upstream to an edge router that acts as the centralized control point and prevents direct subscriber-to-subscriber communication. Ethernet Link Aggregation (LAG) provides the resiliency in this type of network.

  Layer 2 DSLAMs cannot interpret IGMP, so they cannot selectively replicate IPTV channels.

- **Layer–3 aware MSAN**—This IP-aware MSAN can interpret and respond to IGMP requests by locally replicating a multicast stream and forwarding the stream to any subscriber requesting it. Layer 3 awareness is important when supporting IPTV traffic to perform channel changes (sometimes referred to as *channel zaps*). Static IP-aware MSANs always receive all multicast television channels. They do not have the ability to request that specific channels be forwarded to the DSLAM. Dynamic IP-aware DSLAMs, however, can inform the network to begin (or discontinue) sending individual channels to the DSLAM. Configuring IGMP proxy or IGMP snooping on the DSLAM accomplishes this function.

- **Layer–3 MSAN**—These MSANs use IP routing functionality rather than Layer 2 technologies to forward traffic. The advantage of this forwarding method is the ability to support multiple upstream links going to different upstream routers and improving network resiliency. However, to accomplish this level of resiliency, you must assign a separate IP subnetwork to each MSAN, adding a level of complexity that can be more difficult to maintain or manage.

In choosing a MSAN type, refer to Figure 2 on page 18:

Figure 2: Choosing an MSAN Type

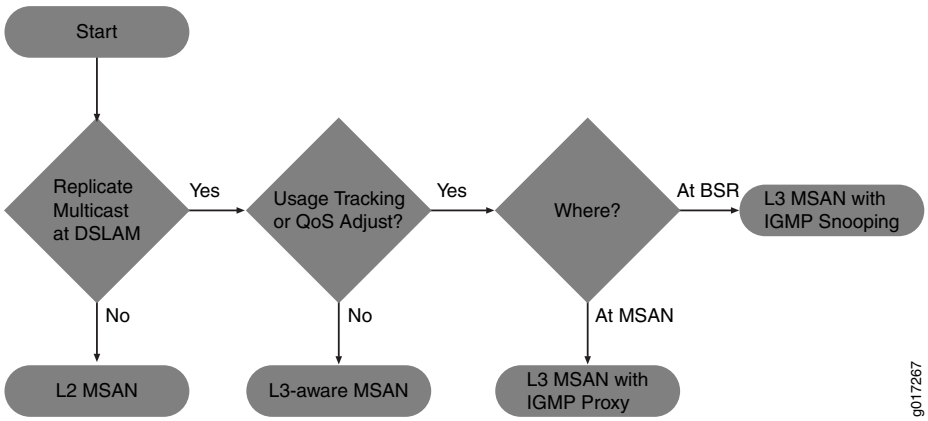

Related Documentation

- Ethernet MSAN Aggregation Options on page 19

Ethernet MSAN Aggregation Options

Each MSAN can connect directly to an edge router (broadband services router or video services router), or an intermediate device (for example, an Ethernet switch) can aggregate MSAN traffic before being sent to the services router. Table 4 on page 19 lists the possible MSAN aggregation methods and under what conditions they are used.

Table 4: Ethernet MSAN Aggregation Methods

Method	When Used
Direct connection	Each MSAN connects directly to the broadband services router and optional video services router.
Ethernet aggregation switch connection	Each MSAN connects directly to an intermediate Ethernet switch. The switch, in turn, connects to the broadband services router or optional video services router.
Ethernet ring aggregation connection	Each MSAN connects to a ring topology of MSANs. The head-end MSAN (the device closest to the upstream edge router) connects to the broadband services router.

You can use different aggregation methods in different portions of the network. You can also create multiple layers of traffic aggregation within the network. For example, an MSAN can connect to a central office terminal (COT), which, in turn, connects to an Ethernet aggregation switch, or you can create multiple levels of Ethernet aggregation switches prior to connecting to the edge router.

Direct Connection

In the direct connection method, each MSAN has a point-to-point connection to the broadband services router. If an intermediate central office exists, traffic from multiple MSANs can be combined onto a single connection using wave-division multiplexing (WDM). You can also connect the MSAN to a video services router. However, this connection method requires that you use a Layer 3 MSAN that has the ability to determine which link to use when forwarding traffic.

When using the direct connection method, keep the following in mind:

- We recommend this approach when possible to simplify network management.
- Because multiple MSANs are used to connect to the services router, and Layer 3 MSANs generally require a higher equipment cost, this method is rarely used in a multiedge subscriber management model.
- Direct connection is typically used when most MSAN links are utilized less than 33 percent and there is little value in combining traffic from multiple MSANs.

Ethernet Aggregation Switch Connection

An Ethernet aggregation switch aggregates traffic from multiple downstream MSANs into a single connection to the services router (broadband services router or optional video services router).

When using the Ethernet aggregation switch connection method, keep the following in mind:

- Ethernet aggregation is typically used when most MSAN links are utilized over 33 percent or to aggregate traffic from lower speed MSANs (for example, 1 Gbps) to a higher speed connection to the services router (for example, 10 Gbps).
- You can use an MX Series router as an Ethernet aggregation switch. For information about configuring the MX Series router in Layer 2 scenarios, see the *Junos OS Layer 2 Switching and Bridging Library for Routing Devices* or the *Ethernet Networking Feature Guide for MX Series Routers*.

Ring Aggregation Connection

In a ring topology, the remote MSAN that connects to subscribers is called the remote terminal (RT). This device can be located in the outside plant (OSP) or in a remote central office (CO). Traffic traverses the ring until it reaches the central office terminal (COT) at the head-end of the ring. The COT then connects directly to the services router (broadband services router or video services router).

NOTE: **The RT and COT must support the same ring resiliency protocol.**

You can use an MX Series router in an Ethernet ring aggregation topology. For information about configuring the MX Series router in Layer 2 scenarios, see the *Junos OS Layer 2 Switching and Bridging Library for Routing Devices* or the *Ethernet Networking Feature Guide for MX Series Routers*.

Related Documentation

- Multiservice Access Node Overview on page 17

CHAPTER 4

Broadband Subscriber Management Solution Software Overview

Broadband Subscriber Management VLAN Architecture Overview

The subscriber management logical network architecture is as important as the physical network architecture. You configure the logical portion of the subscriber management network using virtual local area networks (VLANs).

Three VLAN models deliver multiple services to subscribers. These models include the following:

- **Service VLAN**—The service VLAN (S-VLAN) provides many-to-one (N:1) subscriber-to-service connectivity: The service VLAN carries a service (for example, data, video, or voice) to all subscribers instead of having different services share a VLAN. Adding a new service requires adding a new VLAN and allocating bandwidth to the new service. The service VLAN model enables different groups that are using the broadband network (for example, external application providers) to manage a given service. One limitation of service VLANs is the absence of any logical isolation between user sessions at the VLAN level. This lack of isolation requires that the multiservice access node (MSAN) and broadband services router provide the necessary security filtering.
- **Customer VLAN**—The customer VLAN (C-VLAN) provides one-to-one (1:1) subscriber-to-service connectivity: One VLAN carries all traffic to each subscriber on the network. Having a single VLAN per subscriber simplifies operations by providing a 1:1 mapping of technology (VLANs) to subscribers. You can also understand what applications any subscriber is using at any given time. Because you use only one VLAN to carry traffic to each subscriber, this approach is not affected when adding new services. However, using a pure C-VLAN model consumes more bandwidth because

a single television channel being viewed by multiple subscribers is carried across the network several times—once on each C-VLAN. This approach requires a more scalable, robust edge router that can support several thousand VLANs.

- **Hybrid C-VLAN**—The hybrid VLAN combines the best of both previous VLANs by using one VLAN per subscriber to carry unicast traffic and one shared multicast VLAN (M-VLAN) for carrying broadcast (multicast) television traffic. You can use both the *pure* and *hybrid* C-VLAN models in different portions of the network, depending upon available bandwidth and MSAN capabilities.

NOTE: The term *C-VLAN*, when used casually, often refers to a *hybrid* C-VLAN implementation.

We recommend using one of the C-VLAN models to simplify configuration and management when expanding services. However, some MSANs are limited to the number of VLANs they can support, limiting the ability to use either C-VLAN model.

NOTE: Most MSANs can support the service VLAN model.

Broadband Subscriber Management VLANs Across an MSAN

You configure VLANs to operate between the MSAN and the edge router (broadband services router or video services router). However, the MSAN might modify VLAN identifiers before forwarding information to the subscriber in the following ways:

NOTE: Not all MSANs support these options.

- The VLAN identifiers can be carried within the ATM VCs or they can be removed. The value of keeping the VLAN header is that it carries the IEEE 802.1p Ethernet priority bits. These priority bits can be added to upstream traffic by the residential gateway, allowing the DSLAM to easily identify and prioritize more important traffic (for example, control and VoIP traffic). Typically, a VLAN identifier of zero (0) is used for this purpose.
- In a C-VLAN model, the MSAN might modify the VLAN identifier so that the same VLAN is sent to each subscriber. This enables the use of the same digital subscriber line (DSL) modem and residential gateway configuration for all subscribers without the need to define a different VLAN for each device.

Customer VLANs and Ethernet Aggregation

The 12-bit VLAN identifier (VLAN ID) can support up to 4095 subscribers. When using an aggregation switch with a C-VLAN topology, and fewer than 4095 subscribers are connected to a single edge router port, the aggregation switch can transparently pass all VLANs. However, if the VLAN can exceed 4095 subscribers per broadband services router port, you must use VLAN stacking (IEEE 802.1ad, also known as Q-in-Q). VLAN stacking includes two VLAN tags—an outer tag to identify the destination MSAN and an

inner tag to identify the subscriber. For downstream traffic (that is, from the broadband services router or Ethernet switch to the MSAN), the outer tag determines which port to forward traffic. The forwarding device then uses the VLAN pop function on this tag before forwarding the traffic with a single tag. The reverse process occurs for upstream traffic.

VLAN stacking is not necessary for S-VLANs or M-VLANs. However, for the hybrid (C-VLAN and M-VLAN) model, the Ethernet switch or services router must be able to pop or push tags onto C-VLAN traffic while not modifying M-VLAN packets.

VLANs and Residential Gateways

One function provided by a residential gateway is to enable each subscriber to have a private (in-home) network, unseen by other broadband subscribers, while enabling the subscriber to have multiple devices connected to the broadband network. This private network is made possible by using Network Address Translation (NAT).

Most conditional access systems (for example, video on demand) require detecting the real IP address of the set-top box (STB). This security measure means that traffic to and from the STB must be bridged, not routed, across all network elements including aggregation switches, MSANs, and residential gateways. NAT cannot be used at the residential gateway for traffic to and from the STB. In addition, some residential gateways associate VLANs (or ATM virtual circuits) with ports. Traffic on a given VLAN is always forwarded to specific downstream port. Use caution when mapping VLANs on an MSAN.

Related Documentation

- *Static Subscriber Interfaces and VLAN Overview* in the *Junos OS Subscriber Management and Services Library*.

Broadband Subscriber Management IGMP Model Overview

In an IPTV network, channel changes occur when a set-top box (STB) sends IGMP commands that inform an upstream device (for example, a multiservice access node [MSAN] or services router) whether to start or stop sending multicast groups to the subscriber. In addition, IGMP hosts periodically request notification from the STB about which channels (multicast groups) are being received.

You can implement IGMP in the subscriber management network in the following ways:

- **Static IGMP**—All multicast channels are sent to the MSAN. When the MSAN receives an IGMP request to start or stop sending a channel, it adds the subscriber to the multicast group and then discards the IGMP packet.
- **IGMP Proxy**—Only multicast channels currently being viewed are sent to the MSAN. If the MSAN receives a request to view a channel that is not currently being forwarded to the MSAN, it forwards the request upstream. However, the upstream device does not see all channel change requests from each subscriber, limiting bandwidth control options.

- **IGMP Snooping**—Only multicast channels currently being viewed are sent to the MSAN. The MSAN forwards all IGMP requests upstream, unaltered, even if it is already receiving the channel. The upstream device sees all channel change requests from each subscriber. Using IGMP snooping enables the broadband services router to determine the mix of services and the bandwidth requirements of each subscriber and adjust the bandwidth made available to each service.
- **IGMP Passthrough**—The MSAN transparently passes IGMP packets upstream to the broadband services router.

IGMP hosts (sources) also periodically verify that they are sending the correct traffic by requesting that each client send information about what multicast groups it wants to receive. The responses to this *IGMP query* can result in a substantial upstream traffic burst.

IGMPv2 is the minimum level required to support IPTV, and is the most widely deployed. Emerging standards specify IGMPv3.

Related Documentation

- *Dynamic IGMP Configuration Overview* in the *Junos OS Subscriber Management and Services Library*.

DHCP and Broadband Subscriber Management Overview

You use DHCP in broadband networks to provide IP address configuration and service provisioning. DHCP, historically a popular protocol in LANs, works well with Ethernet connectivity and is becoming increasingly popular in broadband networks as a simple, scalable solution for assigning IP addresses to subscriber home PCs, set-top boxes (STBs), and other devices.

The Junos OS broadband subscriber management solution currently supports the following DHCP allocation models:

- DHCP Local Server
- DHCP Relay

DHCP uses address assignment pools from which to allocate subscriber addresses. Address-assignment pools support both dynamic and static address assignment:

- Dynamic address assignment—A subscriber is automatically assigned an address from the address-assignment pool.
- Static address assignment—Addresses are reserved and always used by a particular subscriber.

NOTE: **Addresses that are reserved for static assignment are removed from the dynamic address pool and cannot be assigned to other clients.**

Extended DHCP Local Server and Broadband Subscriber Management Overview

You can enable the services router to function as an extended DHCP local server. As an extended DHCP local server the services router, and not an external DHCP server, provides an IP address and other configuration information in response to a client request. The extended DHCP local server supports the use of external AAA authentication services, such as RADIUS, to authenticate DHCP clients.

Extended DHCP Relay and Broadband Subscriber Management Overview

You can configure extended DHCP relay options on the router and enable the router to function as a DHCP relay agent. A DHCP relay agent forwards DHCP request and reply packets between a DHCP client and a DHCP server. You can use DHCP relay in carrier edge applications such as video and IPTV to obtain configuration parameters, including an IP address, for your subscribers. The extended DHCP relay agent supports the use of external AAA authentication services, such as RADIUS, to authenticate DHCP clients.

Related Documentation

- *Extended DHCP Local Server Overview* in the *Junos OS Subscriber Management and Services Library*.
- *Extended DHCP Relay Agent Overview* in the *Junos OS Subscriber Management and Services Library*.
- *Address-Assignment Pools Overview* in the *Junos OS Subscriber Management and Services Library*.

AAA Service Framework and Broadband Subscriber Management Overview

You use AAA Service Framework for all authentication, authorization, accounting, address assignment, and dynamic request services that the services router uses for network access. The framework supports authentication and authorization through external servers, such as RADIUS. The framework also supports accounting and dynamic-request CoA and disconnect operations through external servers, and address assignment through a combination of local address-assignment pools and RADIUS.

NOTE: **The broadband subscriber management solution currently supports the use of only RADIUS servers.**

The broadband services router interacts with external servers to determine how individual subscribers access the broadband network. The router also obtains information from external servers for the following:

- Methods used for authentication and accounting.
- How accounting statistics are collected and used.
- How dynamic requests are handled.

Related Documentation

- *AAA Service Framework Overview* in the *Junos OS Subscriber Management and Services Library*.
- *RADIUS-Initiated Change of Authorization (CoA) Overview* in the *Junos OS Subscriber Management and Services Library*.
- *RADIUS-Initiated Disconnect Overview* in the *Junos OS Subscriber Management and Services Library*.

Class of Service and Broadband Subscriber Management Overview

Class of service (CoS) is a mechanism that enables you to divide traffic into classes and offer various levels of throughput and acceptable packet loss when congestion occurs. CoS also provides the option of using differentiated services when best-effort traffic delivery is insufficient. You can also configure the services router to provide hierarchical scheduling for subscribers by dynamically adding or deleting queues when subscribers require services.

By using a dynamic profile, you can provide all subscribers in your network with default CoS parameters when they log in. For example, you can configure an access dynamic profile to specify that all subscribers receive a basic data service. If you use RADIUS variables in the dynamic profile, you can enable the service to be activated for those subscribers at login. You can also use variables to configure a service profile that enables subscribers to activate a service or upgrade to different services through RADIUS change-of-authorization (CoA) messages following initial login.

Related Documentation

- *CoS for Subscriber Access Overview* in the *Junos OS Subscriber Management and Services Library*.

Policy and Control for Broadband Subscriber Management Overview

You can use the Juniper Networks Session and Resource Control (SRC) software to implement policy and control in the subscriber management network. The SRC software provides policy management, subscriber management, and network resource control functions that enable the creation and delivery of services across the network.

For additional information about the Juniper Networks SRC software, go to http://www.juniper.net/techpubs/software/management/src/.

Related Documentation

- Broadband Subscriber Management Overview on page 3
- *Juniper Networks Session and Resource Control (SRC) and JSRC Overview*

PART 2

Configuration

CHAPTER 5

Broadband Subscriber Management Triple Play Configuration

Broadband Subscriber Management Solution Topology and Configuration Elements

The network topology for the broadband subscriber management solution focuses on configuring the access network to which the MX Series routers connect. There are many possible broadband subscriber management configurations. Figure 3 on page 30 illustrates an example of a basic DHCP topology model.

Figure 3: Basic Subscriber Management Solution Topology for a DHCP Subscriber Network

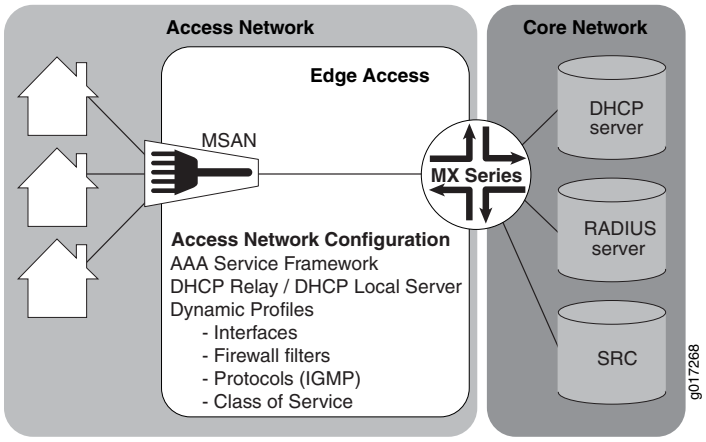

When configuring the broadband subscriber management solution, specific configuration elements come into play. In one form or another, you must configure each of these elements for the subscriber management solution to function.

The configuration elements include the following:

- Subscriber network VLAN configuration
- AAA Service Framework configuration
- Addressing server or addressing server access configuration
- Dynamic profile configuration
- Core network configuration

Related Documentation

- Triple Play Subscriber Management Network Topology Overview on page 31
- Configuring Top-Level Broadband Subscriber Management Elements on page 31

Subscriber Management Licensing

To enable some Junos OS subscriber management software features or router scaling levels, you must purchase, install, and manage certain software license packs. The presence on the router of the appropriate software license keys (passwords) determines whether you can configure and use certain features or configure a feature to a predetermined scale.

For information about how to purchase Juniper Networks Junos OS licenses, contact your Juniper Networks sales representative. For information about installing and managing software licenses that pertain to your broadband subscriber management network, see the *Installation and Upgrade Guide*.

Related Documentation

- Configuring Top-Level Broadband Subscriber Management Elements on page 31

Triple Play Subscriber Management Network Topology Overview

This configuration explains the basics in configuring a basic triple-play (data, voice, and video) network. Figure 4 on page 31 provides the reference topology for this configuration example.

Figure 4: Triple Play Network Reference Topology

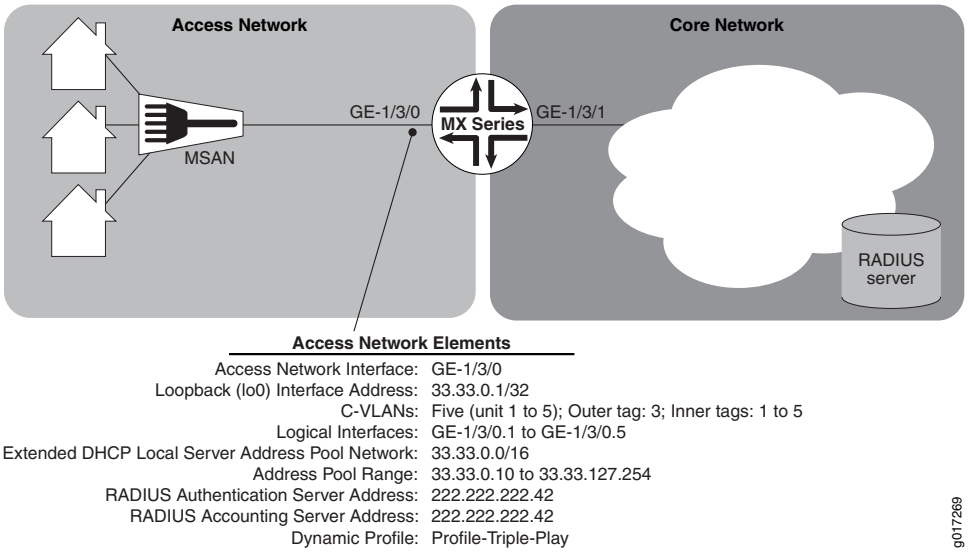

Related Documentation

- Configuring Top-Level Broadband Subscriber Management Elements on page 31

Configuring Top-Level Broadband Subscriber Management Elements

When configuring an MX Series router to act as a broadband services router (BSR) or video services router (VSR), you initially define elements that the router uses to define both subscriber access and the level of service a subscriber can have in your network. Many of these elements are profiles (groups of configuration statements) or static configuration components (like firewall filters) that typically do not change after you create them. After you define these elements, the router can use them to enable subscribers to gain access to your network.

The top-level steps for configuring the edge access in the subscriber management network include the following:

1. Configure the subscriber loopback interface and VLANs.

   See "Configuring Static Customer VLANs for the Broadband Subscriber Management Solution" on page 33.

2. Configure a class of service profile.

   See "Configuring a Global Class of Service Profile for the Broadband Subscriber Management Solution" on page 36.

3. Configure a firewall filter for use with the dynamic profile.

   See "Configuring Dynamic Firewall Filter Services for Use in Dynamic Profiles" on page 42.

4. Configure AAA Framework Services.

   See "Configuring AAA Service Framework for the Broadband Subscriber Management Solution" on page 43.

5. Configure an address assignment pool for use by the address server.

   See "Configuring Address Server Elements for the Broadband Subscriber Management Solution" on page 45.

6. Configure DHCP local server to assign subscriber addresses.

   See "Configuring Address Server Elements for the Broadband Subscriber Management Solution" on page 45.

Related Documentation

- Triple Play Subscriber Management Network Topology Overview on page 31
- Broadband Subscriber Management Solution Topology and Configuration Elements on page 29

Configuring a Loopback Interface for the Broadband Subscriber Management Solution

You must configure a loopback interface for use in the subscriber management access network. The loopback interface is automatically used for unnumbered interfaces.

NOTE: If you do not configure the loopback interface, the routing platform 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.

To configure a loopback interface:

1. Edit the loopback interface.

   ```
   [edit]
   user@host# edit interfaces lo0
   ```

2. Edit the loopback interface unit.

   ```
   [edit interfaces lo0]
   user@host# edit unit 0
   ```

3. Edit the loopback interface family.

   ```
   [edit interfaces lo0 unit 0]
   user@host# edit family inet
   ```

4. Specify the loopback interface address.

```
[edit interfaces lo0 unit 0]
user@host# set address 33.33.0.1/32
```

Related Documentation

- Configuring Top-Level Broadband Subscriber Management Elements on page 31
- *Junos OS Network Interfaces Library for Routing Devices*

Configuring Static Customer VLANs for the Broadband Subscriber Management Solution

In this example configuration, the access interface (**ge-1/3/0**) connects to a device (that is, a DSLAM) on the access side of the network. You can define static customer VLANs (C-VLANs) for use by the access network subscribers.

For a PPPoE solution, to configure the customer VLANs:

1. Edit the access side interface.

   ```
   [edit]
   user@host# edit interfaces ge-1/3/0
   ```

2. Edit the interface unit for the first VLAN.

   ```
   [edit interfaces ge-1/3/0]
   user@host# edit unit 1
   ```

3. Define the VLAN tags for the first VLAN.

   ```
   [edit interfaces ge-1/3/0 unit 1]
   user@host# set vlan-tags outer 3 inner 1
   ```

4. Repeat steps 2 through 3for VLAN interface units 2 through 5.

For a DHCP solution, to configure the customer VLANs:

1. Edit the access side interface.

   ```
   [edit]
   user@host# edit interfaces ge-1/3/0
   ```

2. Edit the interface unit for the first VLAN.

   ```
   [edit interfaces ge-1/3/0]
   user@host# edit unit 1
   ```

3. Define the VLAN tags for the first VLAN.

   ```
   [edit interfaces ge-1/3/0 unit 1]
   user@host# set vlan-tags outer 3 inner 1
   ```

4. Specify that you want to create IPv4 demux interfaces.

   ```
   [edit interfaces ge-1/3/0 unit 1]
   user@host# set demux-source inet
   ```

5. Edit the family for the first VLAN.

   ```
   [edit interfaces ge-1/3/0 unit 1]
   user@host# edit family inet
   ```

6. Define the unnumbered address and the preferred source address for the first VLAN.

   [edit interfaces ge-1/3/0 unit 1 family inet]
   user@host# **set unnumbered-address lo0.0 preferred-source-address 33.33.0.1**

7. Repeat steps 2 through 6for VLAN interface units 2 through 5.

Related Documentation

- Configuring Top-Level Broadband Subscriber Management Elements on page 31
- *Junos OS Network Interfaces Library for Routing Devices*

Configuring Dynamic Customer VLANs for the Broadband Subscriber Management Solution

In this example configuration, the access interface (**ge-1/3/0**) connects to a device (that is, a DSLAM) on the access side of the network. This procedure enables the dynamic creation of up to five customer VLANs (C-VLANs) for use by the access network subscribers.

To configure dynamic VLANs for the solution:

1. Configure a dynamic profile for dynamic VLAN creation.

   a. Name the profile.

      [edit]
      user@host# **edit dynamic-profiles VLAN-PROF**

   b. Define the interfaces statement with the internal **$junos-interface-ifd-name** variable used by the router to match the interface name of the receiving interface.

      [edit dynamic-profiles VLAN-PROF]
      user@host# **edit interfaces $junos-interface-ifd-name**

   c. Define the unit statement with the predefined **$junos-interface-unit** variable:

      [edit dynamic-profiles VLAN-PROF interfaces “$junos-interface-ifd-name”]
      user@host# **set unit $junos-interface-unit**

   d. (Optional) To configure the router to respond to any ARP request, specify the proxy-arp statement.

      [edit dynamic-profiles VLAN-PROF interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit"]
      user@host# **set proxy-arp**

   e. Specify that you want to create IPv4 demux interfaces.

      [edit dynamic-profiles VLAN-PROF interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit"]
      user@host# **set demux-source inet**

   f. Specify the VLAN ID variable.

      [edit dynamic-profiles VLAN-PROF interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit"]
      user@host# **set vlan-tags outer $junos-stacked-vlan-id**

The variable is dynamically replaced with an outer VLAN ID within the VLAN range specified at the **[edit interfaces]** hierarchy level.

g. Specify the inner VLAN ID variable.

```
[edit dynamic-profiles VLAN-PROF interfaces "$junos-interface-ifd-name" unit
   "$junos-interface-unit"]
user@host# set vlan-tags inner $junos-vlan-id
```

The variable is dynamically replaced with an inner VLAN ID within the VLAN range specified at the **[edit interfaces]** hierarchy level.

h. Specify the family type.

```
[edit dynamic-profiles VLAN-PROF interfaces "$junos-interface-ifd-name" unit
   "$junos-interface-unit"]
user@host# set family inet
```

i. (Optional) Enable IP and MAC address validation for dynamic IP demux interfaces in a dynamic profile.

```
[edit dynamic-profiles VLAN-PROF interfaces "$junos-interface-ifd-name" unit
   "$junos-interface-unit" family inet]
user@host# set mac-validate strict
```

j. Specify the unnumbered address and preferred source address.

```
[edit dynamic-profiles VLAN-PROF interfaces "$junos-interface-ifd-name" unit
   "$junos-interface-unit" family inet]
user@host# set unnumbered-address lo.0 preferred-source-address 33.33.0.1
```

2. Associate the dynamic profile with the VLAN interface.

a. Access the interface that you want to use for creating VLANs.

```
[edit interfaces]
user@host# edit interfaces ge-1/3/0
```

b. Specify that you want to automatically configure VLAN interfaces.

```
[edit interfaces ge-1/3/0]
user@host# edit auto-configure
```

c. Specify that you want to configure stacked VLANs.

```
[edit interfaces ge-1/3/0 auto-configure]
user@host# edit stacked-vlan-ranges
```

d. Specify the dynamic VLAN profile that you want the interface to use.

```
[edit interfaces ge-1/3/0 auto-configure stacked-vlan-ranges]
user@host# set dynamic-profile VLAN-PROF
```

3. Specify the Ethernet packet type that the VLAN dynamic profile can accept.

```
[edit interfaces ge-1/3/0 auto-configure stacked-vlan-ranges dynamic-profile
   VLAN-PROF]
user@host# set accept inet
```

4. Define VLAN ranges for use by the dynamic profile when dynamically creating VLAN IDs. For this solution, specify the outer and inner stacked VLAN ranges that you want the dynamic profile to use. To mimic the static VLAN configuration, the following example specifies an outer stacked VLAN ID range of **3–3** (enabling only the outer range of 3) and an inner stacked VLAN ID range of **1–5** (enabling a range from 1 through 5 for the inner stacked VLAN ID).

```
[edit interfaces ge-1/3/0 auto-configure stacked-vlan-ranges dynamic-profile
  VLAN-PROF]
user@host# set ranges 3–3,1–5
```

Related Documentation

- Configuring Top-Level Broadband Subscriber Management Elements on page 31
- Broadband Subscriber Management VLAN Architecture Overview on page 21
- *Dynamic 802.1Q VLAN Overview* in the *Junos OS Network Interfaces Library for Routing Devices*
- *Configuring VLAN Dynamic Profiles* in the *Junos OS Subscriber Management and Services Library*
- *Configuring VLAN Interfaces to Use Dynamic Profiles* in the *Junos OS Subscriber Management and Services Library*
- *Configuring Which VLAN Ethernet Packet Types Dynamic Profiles Can Accept* in the *Junos OS Subscriber Management and Services Library*
- *Configuring VLAN Ranges for Use with Dynamic Profiles* in the *Junos OS Subscriber Management and Services Library*
- *Junos OS Network Interfaces Library for Routing Devices*

Configuring a Global Class of Service Profile for the Broadband Subscriber Management Solution

Junos OS CoS enables you to divide traffic into classes and offer various levels of throughput and packet loss (when congestion occurs) in accordance to service rules that you specify. The Junos OS CoS features provide a set of mechanisms that you can use to provide differentiated (video, voice, and data) services over the same network for subscribers.

- Configuring a Class of Service Profile on page 37
- Configuring CoS Fowarding Classes on page 37
- Configuring CoS Schedulers on page 38
- Configuring Scheduler Maps on page 39
- Configuring CoS Classifiers on page 40
- Configuring CoS Interface Properties on page 41

Configuring a Class of Service Profile

You can configure class of service (CoS) for all subscribers that successfully establish connection to the broadband network. After you create the CoS profile, you can attach it to subscriber interfaces using a dynamic profile.

Configuring a CoS profile includes the following general steps:

1. Configuring forwarding classes.
2. Configuring schedulers.
3. Configuring scheduler maps.
4. Configuring classifiers.
5. Configuring CoS interface properties.

In the configuration we build in this section, we configure three forwarding classes, each with its own scheduler, and an IP precedence classifier for the traffic destined for the access network. Table 5 on page 37 provides an overview of the queue configuration:

Table 5: Class of Service Queue Configuration

Differentiated Services Classification	Bandwidth	Priority	Purpose
Expedited forwarding (EF)	128 Kbps	strict high	voice traffic
Assured forwarding (AF)	29.4 Mbps	low	video traffic
Best effort (BE)	remainder	low	data traffic

NOTE: **The network control forwarding class is not configured in this solution.**

Configuring CoS Fowarding Classes

Forwarding classes identify output queues for packets. For a classifier to assign an output queue to each packet, it must associate the packet with one of the following forwarding classes:

- Expedited forwarding (EF)—Provides a low loss, low latency, low jitter, assured bandwidth, end-to-end service.
- Assured forwarding (AF)—Provides a group of values you can define and includes four subclasses: AF1, AF2, AF3, and AF4, each with three drop probabilities: low, medium, and high.

- Best effort (BE)—Provides no service profile. For the BE forwarding class, loss priority is typically not carried in a class-of-service (CoS) value, and random early detection (RED) drop profiles are more aggressive.
- Network control (NC)—This class is typically high priority because it supports protocol control.

NOTE: The MX Series router enables you to configure up to eight forwarding class queues.

To configure forwarding class queues:

1. Edit the best effort queue.

   ```
   [edit]
   user@host# edit class-of-service forwarding-classes queue 0
   ```

2. Name the queue.

   ```
   [edit class-of-service forwarding-classes queue 0]
   user@host# set fc_be
   ```

3. Edit the expedited forwarding queue.

   ```
   [edit]
   user@host# edit class-of-service forwarding-classes queue 1
   ```

4. Name the queue.

   ```
   [edit class-of-service forwarding-classes queue 1]
   user@host# set fc_ef
   ```

5. Edit the assured forwarding queue.

   ```
   [edit]
   user@host# edit class-of-service forwarding-classes queue 2
   ```

6. Name the queue.

   ```
   [edit class-of-service forwarding-classes queue 1]
   user@host# set fc_af
   ```

Configuring CoS Schedulers

CoS schedulers define the properties of output queues. These properties can include the amount of interface bandwidth assigned to the queue, the size of the memory buffer allocated for storing packets, the priority of the queue, and the random early detection (RED) drop profiles associated with the queue.

To configure CoS schedulers for the existing queues:

1. Create a scheduler and name it for the best effort traffic.

   ```
   [edit]
   user@host# edit class-of-service schedulers sched_be
   ```

2. Define the best effort scheduler buffer size.

```
[edit class-of-service schedulers sched_be]
user@host# set buffer-size remainder
```

3. Set the priority of the best effort scheduler.

```
[edit class-of-service schedulers sched_be]
user@host# set priority low
```

4. Create a scheduler and name it for the expedited forwarding traffic.

```
[edit]
user@host# edit class-of-service schedulers sched_ef
```

5. Configure the transmit rate for the expedited forwarding scheduler.

```
[edit class-of-service schedulers sched_ef]
user@host# set transmit-rate 128k
```

6. Define the expedited forwarding scheduler buffer size.

```
[edit class-of-service schedulers sched_ef]
user@host# set buffer-size remainder
```

7. Set the priority of the expedited forwarding scheduler.

```
[edit class-of-service schedulers sched_ef]
user@host# set priority strict-high
```

8. Create a scheduler and name it for the assured forwarding traffic.

```
[edit]
user@host# edit class-of-service schedulers sched_af
```

9. Configure the transmit rate for the assured forwarding scheduler.

```
[edit class-of-service schedulers sched_af]
user@host# set transmit-rate 29400000
```

10. Define the assured forwarding scheduler buffer size.

```
[edit class-of-service schedulers sched_af]
user@host# set buffer-size remainder
```

11. Set the priority of the expedited forwarding scheduler.

```
[edit class-of-service schedulers sched_af]
user@host# set priority low
```

Configuring Scheduler Maps

After configuring both CoS forwarding classes and schedulers, you must use scheduler maps to associate them.

To map CoS forwarding classes to schedulers:

1. Create a forwarding map and name it.

```
[edit]
user@host# edit class-of-service scheduler-maps SchedulerMap_Triple_Play_Basic
```

2. Edit the best effort forwarding class queue.

```
[edit class-of-service scheduler-maps SchedulerMap_Triple_Play_Basic]
```

```
user@host#  edit forwarding-class fc_be
```

3. Associate the scheduler that you want this forwarding class to use.

```
[edit class-of-service scheduler-maps SchedulerMap_Triple_Play_Basic forwarding-class
  fc_be]
user@host# set scheduler sched_be
```

4. Edit the expedited forwarding class queue.

```
[edit class-of-service scheduler-maps SchedulerMap_Triple_Play_Basic]
user@host#  edit forwarding-class fc_ef
```

5. Associate the scheduler that you want this forwarding class to use.

```
[edit class-of-service scheduler-maps SchedulerMap_Triple_Play_Basic forwarding-class
  fc_ef]
user@host# set scheduler sched_ef
```

6. Edit the assured forwarding class queue.

```
[edit class-of-service scheduler-maps SchedulerMap_Triple_Play_Basic]
user@host#  edit forwarding-class fc_af
```

7. Associate the scheduler that you want this forwarding class to use.

```
[edit class-of-service scheduler-maps SchedulerMap_Triple_Play_Basic forwarding-class
  fc_af]
user@host# set scheduler sched_af
```

Configuring CoS Classifiers

You can override the default IP precedence classifier by defining a custom classifier. You can then apply the classifier to a logical interface.

To define a custom CoS classifier:

1. Create a Differentiated Services code point (DSCP) classifier and name it.

```
[edit]
user@host# edit class-of-service classifiers dscp Class_DSCP
```

NOTE: DSCP classifiers handle incoming IPv4 packets.

2. Edit the best effort forwarding class queue.

```
[edit class-of-service classifiers dscp Class_DSCP]
user@host# edit forwarding-class fc_be
```

3. Edit the loss priority level for the forwarding class queue.

```
[edit class-of-service classifiers dscp Class_DSCP forwarding-class fc_be]
user@host# edit loss-priority high
```

4. Set code points for the loss priority level.

```
[edit class-of-service classifiers dscp Class_DSCP forwarding-class fc_be loss-priority
  low]
user@host# set code-points be
```

5. Edit the expedited forwarding class queue.

   [edit class-of-service classifiers dscp Class_DSCP]
   user@host# edit forwarding-class fc_ef

6. Edit the loss priority level for the forwarding class queue.

   [edit class-of-service classifiers dscp Class_DSCP forwarding-class fc_ef]
   user@host# edit loss-priority low

7. Set code points for the loss priority level.

   [edit class-of-service classifiers dscp Class_DSCP forwarding-class fc_ef loss-priority low]
   user@host# set code-points ef

8. Edit the assured forwarding class queue.

   [edit class-of-service classifiers dscp Class_DSCP]
   user@host# edit forwarding-class fc_af

9. Edit the loss priority level for the forwarding class queue.

   [edit class-of-service classifiers dscp Class_DSCP forwarding-class fc_af]
   user@host# edit loss-priority low

10. Set code points for the loss priority level.

    [edit class-of-service classifiers dscp Class_DSCP forwarding-class fc_af loss-priority low]
    user@host# set code-points af41

Configuring CoS Interface Properties

Configuring CoS interface properties enables the router to throttle and classify the traffic from the Internet that is sent to subscriber local loops. Limiting the traffic to the access network ensures that the traffic sent to the subscriber local loops does not exceed the current data transmission rate of those lines. Limiting traffic also ensures that changes to subscriber local loop speeds do not cause bandwidth contention at the subscriber's residential gateway. You apply the classifier to the core-facing interface to classify incoming traffic for the queues you are using in the access network.

To configure CoS interfaces:

1. Edit the core CoS interface you want to configure.

   [edit]
   user@host# edit class-of-service interfaces ge-1/3/0

2. Edit the interface shaping rate.

   [edit class-of-service interfaces ge-1/3/0]
   user@host# edit class-of-service interfaces ge-1/3/0 shaping-rate

3. Set the shaping rate value to throttle traffic to the subscriber local loops.

   [edit class-of-service interfaces ge-1/3/0 shaping-rate]
   user@host# set 500m

4. Edit the interface connected to the core network.

```
[edit]
user@host# edit class-of-service interfaces ge-1/3/1
```

5. Edit the interface unit.

```
[edit class-of-service interfaces ge-1/3/1]
user@host# edit unit 0
```

6. Edit the interface unit classifiers.

```
[edit class-of-service interfaces ge-1/3/1 unit 0]
user@host# edit classifiers
```

7. Apply the classifier to the interface to classify traffic coming from the Internet.

```
[edit class-of-service interfaces ge-1/3/1 unit 0 classifiers]
user@host# set dscp Class_DSCP
```

Configuring Dynamic Firewall Filter Services for Use in Dynamic Profiles

Firewall filters provide rules that define whether to permit or deny packets that are transiting an interface on a router. You can configure firewall filters for use in dynamic profiles. After you configure dynamic firewall filters, you can specify which filters you want to apply to subscriber interfaces using a dynamic profile.

To create a firewall filter:

1. Create and name a firewall filter.

```
[edit]
user@host# edit firewall filter fw_fltr_af41
```

2. Specify the filter to be interface specific.

```
[edit firewall filter fw_fltr_af41]
user@host# set interface-specific (Firewall Filters)
```

3. Edit a first term for the firewall filter.

```
[edit firewall filter fw_fltr_af41]
user@host# edit term 1
```

4. Set the **from** match condition.

```
[edit firewall filter fw_fltr_af41 term 1]
user@host# set from dscp af41
```

5. Set the **then** action to take when a match occurs.

```
[edit firewall filter fw_fltr_af41 term 1]
user@host# set then count c2 accept
```

6. Edit a second term for the firewall filter.

```
[edit firewall filter fw_fltr_af41]
user@host# edit term 2
```

7. Set the **then** action to take when a match occurs for term 2.

```
[edit firewall filter fw_fltr_af41 term 2]
```

```
user@host# set then accept
```

8. Apply the dynamic firewall filter to interfaces using a dynamic profile.

   See "Configuring a DHCP Dynamic Profile for the Triple Play Solution" on page 49.

Related Documentation

- Configuring Top-Level Broadband Subscriber Management Elements on page 31
- *Dynamic Firewall Filters Overview* in the *Junos OS Subscriber Management and Services Library*.
- *Dynamic Profiles Overview* in the *Junos OS Subscriber Management and Services Library*.
- *Routing Policy Feature Guide for Routing Devices*

Configuring AAA Service Framework for the Broadband Subscriber Management Solution

Configuring RADIUS Server Access Information

Define the RADIUS server address and secret data that RADIUS access profiles can reference. Define an access profile that includes specific RADIUS configuration.

To configure RADIUS server access:

1. Edit router access to the RADIUS server.

   ```
   [edit]
   user@host# edit access radius-server
   ```

2. Set the address to the RADIUS server.

   ```
   [edit access radius-server]
   user@host# set 222.222.222.42
   ```

3. Edit the RADIUS server.

   ```
   [edit access radius-server]
   user@host# edit 222.222.222.42
   ```

4. Configure the source address for the RADIUS server.

   ```
   [edit access radius-server 222.222.222.42]
   user@host# set source-address 222.222.222.1
   ```

5. Configure the secret for the RADIUS server.

   ```
   [edit access radius-server 222.222.222.42]
   user@host# set secret "$EcReTRad1uSdAta4f0rTh3rtR"
   ```

Configuring RADIUS Server Access Profile

You can define a RADIUS access profile that references defined RADIUS servers and includes specific RADIUS configuration for authentication and accounting.

To configure a RADIUS access profile:

1. Create and name a RADIUS access profile.

   ```
   [edit]
   user@host# edit access profile AccessProfile_general
   ```

2. Edit the order in which authentication mechanisms are used.

   ```
   [edit access profile AccessProfile_general]
   user@host# set authentication-order radius
   ```

3. Edit the RADIUS access addresses.

   ```
   [edit access profile AccessProfile_general]
   user@host# edit access profile AccessProfile_general radius
   ```

4. Set the address or address list for the RADIUS authentication server.

   ```
   [edit access profile AccessProfile_general radius]
   user@host# set authentication-server 222.222.222.42
   ```

5. Set the address or address list for the RADIUS accounting server.

   ```
   [edit access profile AccessProfile_general radius]
   user@host# set accounting-server 222.222.222.42
   ```

6. Edit the RADIUS accounting values for the access profile.

   ```
   [edit access profile AccessProfile_general]
   user@host# edit accounting
   ```

7. Set the RADIUS accounting order.

   ```
   [edit access profile AccessProfile_general accounting]
   user@host# set order radius
   ```

8. Specify that RADIUS accounting stop when a user fails authentication but is granted access.

   ```
   [edit access profile AccessProfile_general accounting]
   user@host# set accounting-stop-on-failure
   ```

9. Specify that RADIUS accounting stop when access is denied to a subscriber.

   ```
   [edit access profile AccessProfile_general accounting]
   user@host# set accounting-stop-on-access-deny
   ```

10. Specify that RADIUS provide immediate updates.

    ```
    [edit access profile AccessProfile_general accounting]
    user@host# set immediate-update
    ```

11. Specify the amount of time (in minutes) between RADIUS updates.

    ```
    [edit access profile AccessProfile_general accounting]
    user@host# set update-interval 10
    ```

12. Specify that RADIUS accounting report only subscriber uptime.

    ```
    [edit access profile AccessProfile_general accounting]
    user@host# set statistics time
    ```

Related Documentation

- Configuring Top-Level Broadband Subscriber Management Elements on page 31
- *Configuring RADIUS Server Parameters for Subscriber Access*
- *AAA Service Framework Overview* in the *Junos OS Subscriber Management and Services Library*.

Configuring Address Server Elements for the Broadband Subscriber Management Solution

Configuring a DHCPv4 Address Assignment Pool

Address assignment pools enable you to specify groups of IP addresses that different client applications can share. In this configuration, the extended DHCP local server configuration or the router PPP software uses the address pool to provide addresses to subscribers that are accessing the network.

For PPP, to configure an address assignment pool:

1. Create and name an address assignment pool.

   ```
   [edit]
   user@host# edit access address-assignment pool AddressPool_1
   ```

2. Edit the address pool family.

   ```
   [edit access address-assignment pool AddressPool_1]
   user@host# edit family inet
   ```

3. Define the address pool network address.

   ```
   [edit access address-assignment pool AddressPool_1 family inet]
   user@host# set network 33.33.0.0/16
   ```

4. Set the address range for the network.

   ```
   [edit access address-assignment pool AddressPool_1 family inet]
   user@host# set range all low 33.33.0.10 high 33.33.127.254
   ```

5. Specify which access profile you want to instantiate.

   ```
   [edit]
   user@host# set access-profile AccessProfile_general
   ```

For DHCP local server, to configure an address assignment pool:

1. Create and name an address assignment pool.

   ```
   [edit]
   user@host# edit access address-assignment pool AddressPool_1
   ```

2. Edit the address pool family.

   ```
   [edit access address-assignment pool AddressPool_1]
   user@host# edit family inet
   ```

3. Define the address pool network address.

   [edit access address-assignment pool AddressPool_1 family inet]
   user@host# **set** network **33.33.0.0/16**

4. Set the address range for the network.

   [edit access address-assignment pool AddressPool_1 family inet]
   user@host# **set** range **all low 33.33.0.10 high 33.33.127.254**

5. Edit the family DHCP attributes.

   [edit access address-assignment pool AddressPool_1 family inet]
   user@host# **edit** family **inet dhcp-attributes**

6. Set the maximum lease time.

   [edit access address-assignment pool AddressPool_1 family inet dhcp-attributes]
   user@host# **set** maximum-lease-time **3600**

7. Set the grace period.

   [edit access address-assignment pool AddressPool_1 family inet dhcp-attributes]
   user@host# **set** grace-period **60**

8. Set the router IP address that you want advertised to subscribers.

   [edit access address-assignment pool AddressPool_1 family inet dhcp-attributes]
   user@host# **set** router **33.33.0.1**

9. Specify which access profile you want to instantiate.

   [edit]
   user@host# **set** access-profile **AccessProfile_general**

Configuring Extended DHCP Local Server

You can enable the MX Series router to function as an extended DHCP local server. The extended DHCP local server provides IP addresses and other configuration information to a subscriber logging in to the network.

To configure the DHCP local server:

1. Edit the routing system services.

   [edit]
   user@host# **edit** system services

2. Edit the DHCP local server.

   [edit system services]
   user@host# **edit** dhcp-local-server

3. Define the DHCP pool match order.

   [edit system services dhcp-local-server]
   user@host# **set** pool-match-order **ip-address-first**

4. Set the authentication password.

   [edit system services dhcp-local-server]
   user@host# **set** authentication password **auth-psswrd**

5. Edit the values you want included with the username.

   [edit system services dhcp-local-server]
   user@host# **edit** authentication username-include

6. Set the values you want included with the username.

   [edit system services dhcp-local-server username-include]
   user@host# **set** domain-name **yourcompany.com**
   user@host# **set** user-prefix **user-defined-prefix**

7. Create and name a DHCP local server group.

   [edit system services dhcp-local-server]
   user@host# **edit** group **dhcp-ls-group**

8. Specify a dynamic profile that you want the DHCP local server group to use.

   [edit system services dhcp-local-server group dhcp-ls-group]
   user@host# **set** dynamic-profile **Profile-Triple_Play**

9. Assign interfaces to the group.

   [edit system services dhcp-local-server group dhcp-ls-group]
   user@host# **set** interface **ge-1/3/0.1 upto ge-1/3/0.5**

10. Edit the DHCP local server trace options.

    [edit system processes dhcp-service]
    user@host# **edit interface-traceoptions**

11. Specify a log file into which you want trace option information to be saved.

    [edit system processes dhcp-service interface-traceoptions]
    user@host# **set file dhcp-server-msgs.log**

12. Specify the DHCP local server message operations that you want saved in the log file.

    [edit system processes dhcp-service interface-traceoptions]
    user@host# **set flag all**

Related Documentation

- Configuring Top-Level Broadband Subscriber Management Elements on page 31
- *Address-Assignment Pools Overview* in the *Junos OS Subscriber Management and Services Library*.
- *Extended DHCP Local Server Overview* in the *Junos OS Subscriber Management and Services Library*.

Configuring a PPPoE Dynamic Profile for the Triple Play Solution

A dynamic profile is a set of characteristics, defined in a type of template, that you can use to provide dynamic subscriber access and services for broadband applications. These services are assigned dynamically to interfaces.

NOTE: The following configuration is PPPoE-specific.

To configure a PPPoE dynamic profile:

1. Create and name the dynamic profile.

   ```
   [edit]
   user@host# edit dynamic-profiles Profile-Triple-Play
   ```

2. Edit the profile PPPoE dynamic interface.

   ```
   [edit dynamic-profiles Profile-Triple-Play]
   user@host# edit interfaces pp0
   ```

3. Edit the unit variable.

   ```
   [edit dynamic-profiles Profile-Triple-Play interfaces pp0]
   user@host# edit unit $junos-interface-unit
   ```

4. Edit the PPP options.

   ```
   [edit dynamic-profiles Profile-Triple-Play interfaces pp0 unit "$junos-interface -unit"]
   user@host# edit ppp-options
   ```

5. (Optional) Specify either **chap** or **pap** (or both).

   ```
   [edit dynamic-profiles Profile-Triple-Play interfaces pp0 unit “$junos-interface-unit”
     ppp-options]
   user@host# set chap
   user@host# set pap
   ```

6. Edit the PPPoE options.

   ```
   [edit dynamic-profiles Profile-Triple-Play interfaces pp0 unit “$junos-interface-unit”]
   user@host# edit pppoe-options
   ```

7. Specify the PPPoE underlying interface variable.

   ```
   [edit dynamic-profiles Profile-Triple-Play interfaces pp0 unit “$junos-interface-unit”
     pppoe-options]
   user@host# set underlying-interface $junos-underlying-interface
   ```

8. Define the router to act as a PPPoE server when a PPPoE logical interface is dynamically created.

   ```
   [edit dynamic-profiles Profile-Triple-Play interfaces pp0 unit “$junos-interface-unit”
     pppoe-options]
   user@host# set server
   ```

9. Edit the dynamic interface family.

   ```
   [edit dynamic-profiles Profile-Triple-Play interfaces "$junos-interface-ifd-name" unit
     "$junos-underlying-interface-unit"]
   user@host# edit family inet
   ```

10. Specify the input filter that you want to apply to each dynamic interface when it is created.

    ```
    [edit dynamic-profiles Profile-Triple-Play interfaces "$junos-interface-ifd-name" unit
      "$junos-underlying-interface-unit" family inet]
    user@host# set filter input fltr_af41
    ```

11. Specify the output filter that you want to apply to each dynamic interface when it is created.

```
[edit dynamic-profiles Profile-Triple-Play interfaces "$junos-interface-ifd-name" unit
  "$junos-underlying-interface-unit" family inet]
user@host# set filter output fltr_af41
```

12. Enable the local address to be derived from the specified PPPoE interface (in this case, the loopback address).

```
[edit dynamic-profiles Profile-Triple-Play interfaces "$junos-interface-ifd-name" unit
  "$junos-underlying-interface-unit" family inet]
user@host# set unnumbered-address lo0.0
```

13. Edit dynamic class of service.

```
[edit dynamic-profiles Profile-Triple-Play]
user@host# edit class-of-service
```

14. Edit the dynamic CoS traffic control profile.

```
[edit dynamic-profiles Profile-Triple-Play class-of-service]
user@host# edit traffic-control-profiles
```

15. Create and name a traffic control profile.

```
[edit dynamic-profiles Profile-Triple-Play class-of-service traffic-control-profiles]
user@host# edit TrafficProfile_Triple_Play
```

16. Specify a scheduler map that you want the dynamic CoS traffic control profile to use.

```
[edit dynamic-profiles Profile-Triple-Play class-of-service traffic-control-profiles
  TrafficProfile_Triple_Play]
user@host# set scheduler-map SchedulerMap_Triple_Play_Basic
```

17. Specify the shaping rate that you want the dynamic CoS traffic control profile to use.

```
[edit dynamic-profiles Profile-Triple-Play class-of-service traffic-control-profiles
  TrafficProfile_Triple_Play]
user@host# set shaping-rate 32700000
```

18. Apply CoS to the dynamic interfaces and apply an output traffic control profile.

```
[edit dynamic-profiles Profile-Triple-Play class-of-service]
user@host# set interfaces $junos-interface-ifd-name unit
  $junos-underlying-interface-unit output-traffic-control-profileotcp-profile
```

Related Documentation

- Configuring Top-Level Broadband Subscriber Management Elements on page 31
- *Dynamic Profiles Overview* in the *Junos OS Subscriber Management and Services Library*.

Configuring a DHCP Dynamic Profile for the Triple Play Solution

A dynamic profile is a set of characteristics, defined in a type of template, that you can use to provide dynamic subscriber access and services for broadband applications. These services are assigned dynamically to interfaces.

NOTE: The following configuration is DHCP-specific.

To configure a DHCP dynamic profile:

1. Create and name the dynamic profile.

   ```
   [edit]
   user@host# edit dynamic-profiles Profile-Triple_Play
   ```

2. Edit the profile dynamic interfaces.

   ```
   [edit dynamic-profiles Profile-Triple-Play]
   user@host# edit interfaces
   ```

3. Edit the dynamic interfaces.

   ```
   [edit dynamic-profiles Profile-Triple-Play interfaces]
   user@host# edit $junos-interface-ifd-name unit $junos-underlying-interface-unit
   ```

4. Edit the dynamic interface family.

   ```
   [edit dynamic-profiles Profile-Triple-Play interfaces "$junos-interface-ifd-name" unit
     "$junos-underlying-interface-unit"]
   user@host# edit family inet
   ```

5. Specify the input filter that you want to apply to each dynamic interface when it is created.

   ```
   [edit dynamic-profiles Profile-Triple-Play interfaces "$junos-interface-ifd-name" unit
     "$junos-underlying-interface-unit" family inet]
   user@host# set filter input fltr_af41
   ```

6. Specify the output filter that you want to apply to each dynamic interface when it is created.

   ```
   [edit dynamic-profiles Profile-Triple-Play interfaces "$junos-interface-ifd-name" unit
     "$junos-underlying-interface-unit" family inet]
   user@host# set filter output fltr_af41
   ```

7. Edit dynamic class of service.

   ```
   [edit dynamic-profiles Profile-Triple-Play]
   user@host# edit class-of-service
   ```

8. Edit the dynamic CoS traffic control profile.

   ```
   [edit dynamic-profiles Profile-Triple_Play class-of-service]
   user@host# edit traffic-control-profiles
   ```

9. Create and name a traffic control profile.

   ```
   [edit dynamic-profiles Profile-Triple_Play class-of-service traffic-control-profiles]
   user@host# edit TrafficProfile_Triple_Play
   ```

10. Specify a scheduler map that you want the dynamic CoS traffic control profile to use.

    ```
    [edit dynamic-profiles Profile-Triple_Play class-of-service traffic-control-profile]
    user@host# set scheduler-map SchedulerMap_Triple_Play_Basic
    ```

11. Specify the shaping rate that you want the dynamic CoS traffic control profile to use.

    ```
    [edit dynamic-profiles Profile-Triple_Play class-of-service traffic-control-profile]
    user@host# set shaping-rate 32700000
    ```

12. Apply CoS to the dynamic interfaces and apply an output traffic control profile.

```
[edit dynamic-profiles Profile-Triple_Play class-of-service]
user@host# set interfaces $junos-interface-ifd-name unit
  $junos-underlying-interface-unit output-traffic-control-profile
  TrafficProfile_Triple_Play
```

Related Documentation

- Configuring Top-Level Broadband Subscriber Management Elements on page 31
- *Dynamic Profiles Overview*

CHAPTER 6

Broadband Subscriber Management Triple Play Configuration Statements

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 Trio MPC/MIC interfaces on MX Series Routers.**

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

Related Documentation

- *Configuring the VLAN Ethernet Packet Type for Single-Tag VLAN Dynamic Profiles*
- *Configuring the VLAN Ethernet Packet Type for Stacked VLAN Dynamic Profiles*
- *Configuring VLAN Interfaces for the Layer 2 Wholesale Solution*

access-profile (Routing Instances)

Syntax

access-profile *profile-name*;

Hierarchy Level

[edit],
[edit routing-instances *routing-instances-name*]

Release Information

Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Specify the access profile for use by the master routing instance.

Options

profile-name—Name of the access profile.

Required Privilege Level

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

Related Documentation

- *Configuring Access Components for the DHCP Layer 3 Wholesale Network Solution*
- *Configuring Access Components for the PPPoE Wholesale Network Solution*
- Configuring Address Server Elements for the Broadband Subscriber Management Solution on page 45

accounting (Access Profile)

Syntax

```
accounting {
  accounting-stop-on-access-deny;
  accounting-stop-on-failure;
  address-change-immediate-update;
  coa-immediate-update;
  coa-no-override service-class-attribute;
  duplication;
  duplication-vrf {
    access-profile-name profile-name;
    vrf-name vrf-name;
  }
  immediate-update;
  order [ accounting-method ];
  statistics (time | volume-time);
  update-interval minutes;
  wait-for-acct-on-ack;
}
```

Hierarchy Level

[edit access profile *profile-name*]

Release Information

Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.

Description

Configure RADIUS accounting parameters and enable RADIUS accounting for an access profile.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Configuring Authentication and Accounting Parameters for Subscriber Access*
- *Configuring Per-Subscriber Session Accounting*
- *Understanding RADIUS Accounting Duplicate Reporting*

accounting-server

Syntax accounting-server [*ip-address*];

Hierarchy Level [edit access profile *profile-name* radius]

Release Information Statement introduced in Junos OS Release 9.1.

Description Specify a list of the RADIUS accounting servers used for accounting for DHCP, L2TP, and PPP clients.

Options ***ip-address***—IP version 4 (IPv4) address.

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

Related Documentation
- *Configuring Authentication and Accounting Parameters for Subscriber Access*

accounting-stop-on-access-deny

Syntax accounting-stop-on-access-deny;

Hierarchy Level [edit access profile *profile-name* accounting]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.

Description Configure RADIUS accounting to send an Acct-Stop message when the AAA server refuses a client request for access.

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

Related Documentation
- *Configuring Authentication and Accounting Parameters for Subscriber Access*

accounting-stop-on-failure

Syntax accounting-stop-on-failure;

Hierarchy Level [edit access profile *profile-name* accounting]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.

Description Configure RADIUS accounting to send an Acct-Stop message when client access fails AAA but the AAA server grants access.

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

Related Documentation

- *Configuring Authentication and Accounting Parameters for Subscriber Access*

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;
  (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*],
[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.

The remaining statements are explained separately.

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*
- *negotiate-address*
- *unnumbered-address (Ethernet)*
- *Junos OS Administration Library for Routing Devices*
- *family*

address-assignment (Address-Assignment Pools)

Syntax

```
address-assignment {
  abated-utilization percentage;
  abated-utilization-v6 percentage;
  high-utilization percentage;
  high-utilization-v6 percentage;
  neighbor-discovery-router-advertisement ndra-pool-name;
  pool pool-name {
    family family {
      dhcp-attributes {
        protocol-specific attributes;
      }
      host hostname {
        hardware-address mac-address;
        ip-address ip-address;
      }
      network ip-prefix/<prefix-length>;
      prefix ipv6-prefix;
      range range-name {
        high upper-limit;
        low lower-limit;
        prefix-length prefix-length;
      }
    }
    link pool-name;
  }
}
```

Hierarchy Level [edit access]

Release Information Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description Configure address-assignment pools that can be used by different client applications.

NOTE: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Options *pool-name*—Name assigned to an address-assignment pool.

The remaining statements are explained separately.

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

Related Documentation

- *Address-Assignment Pools Overview*
- *Configuring Address-Assignment Pools*

- *Configuring an Address-Assignment Pool for L2TP LNS with Inline Services*
- *Configuring a DHCP Server on EX Series Switches (CLI Procedure)*

authentication (DHCP Local Server)

Syntax

```
authentication {
  password password-string;
  username-include {
    circuit-type;
    client-id;
    delimiter delimiter-character;
    domain-name domain-name-string;
    interface-name ;
    logical-system-name;
    mac-address;
    option-60;
    option-82 <circuit-id> <remote-id>;
    relay-agent-interface-id;
    relay-agent-remote-id;
    relay-agent-subscriber-id;
    routing-instance-name;
    user-prefix user-prefix-string;
  }
}
```

Hierarchy Level

[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server dhcpv6 group *group-name*],
[edit system services dhcp-local-server group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server ...],
[edit logical-systems *logical-system-name* system services dhcp-local-server ...],
[edit routing-instances *routing-instance-name* system services dhcp-local-server ...]

Release Information

Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure the parameters the router sends to the external AAA server. A group configuration takes precedence over a global DHCP relay or DHCP local server configuration.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Using External AAA Authentication Services with DHCP*

authentication-order

Syntax authentication-order [*authentication-methods*];

Hierarchy Level [edit access profile *profile-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
none option introduced in Junos OS Release 11.2.

Description Set the order in which the Junos OS tries different authentication methods when verifying that a client can access the router or switch. For each login attempt, the software tries the authentication methods in order, from first to last.

Default password

Options *authentication-methods*

- **none**—Grants authentication without examining the client credentials. Can be used, for example, when the Diameter function Gx-Plus is employed for notification during subscriber provisioning.
- **password**—Verify the client using the information configured at the **[edit access profile *profile-name* client *client-name*]** hierarchy level.
- **radius**—Verify the client using RADIUS authentication services.

NOTE: For subscriber access management, you must always specify the radius method. Subscriber access management does not support the password option (the default), and authentication fails when no method is specified.

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

Related Documentation

- *Example: Configuring CHAP Authentication with RADIUS*
- *Specifying the Authentication and Accounting Methods for Subscriber Access*
- *Configuring Access Profiles for L2TP or PPP Parameters*

authentication-server

Syntax authentication-server [*ip-address*];

Hierarchy Level [edit access profile *profile-name* radius]

Release Information Statement introduced in Junos OS Release 9.1.

Description Specify a list of the RADIUS authentication servers used to authenticate DHCP, L2TP, and PPP clients. The servers in the list are also used as RADIUS dynamic-request servers, from which the router accepts and processes RADIUS disconnect requests, CoA requests, and dynamic service activations and deactivations.

Options ***ip-address***—IPv4 address.

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

Related Documentation

- *Configuring RADIUS Server Parameters for Subscriber Access*

auto-configure

Syntax

```
auto-configure {
  vlan-ranges {
    access-profile profile-name;
    authentication {
      password password-string;
      username-include{
        circuit-type;
        delimiter delimiter-character;
        domain-name domain-name-string;
        interface-name;
        mac-address;
        option-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)–(any | high-tag);
    }
    override;
  }
  stacked-vlan-ranges {
    access-profile profile-name;
    authentication {
      password password-string;
      username-include {
        circuit-type;
        delimiter delimiter-character;
        domain-name domain-name-string;
        interface-name;
        mac-address;
        option-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.

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

Related Documentation

- *Configuring VLAN Interfaces to Use Dynamic Profiles*

buffer-size (Schedulers)

Syntax buffer-size (percent *percentage* | remainder | temporal *microseconds*);

Hierarchy Level [edit class-of-service schedulers *scheduler-name*]

Release Information Statement introduced before Junos OS Release 7.4.
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 Routers.

Description Specify buffer size.

NOTE: **On PTX Series Packet Transport Routers, buffer-size cannot be configured on rate-limited queues.**

Default If you do not include this statement, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent, respectively.

Options **percent *percentage***—Buffer size as a percentage of the total buffer.

Range: 0 through 100

remainder—Remaining buffer available.

temporal *microseconds*—Buffer size as a temporal value. The queuing algorithm starts dropping packets when it queues more than a computed number of bytes. This maximum is computed by multiplying the logical interface speed by the configured temporal value.

Range: The ranges vary by platform. See *Buffer Size Temporal Value Ranges by Router Type*.

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

Related Documentation

- *Configuring the Scheduler Buffer Size*
- *Example: Configuring CoS for a PBB Network on MX Series Routers*

chap (Dynamic PPP)

Syntax

```
chap {
   challenge-length minimum minimum-length maximum maximum-length;
}
```

Hierarchy Level

[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 in Junos OS Release 9.5.
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 12.2.

Description

Specify CHAP authentication in a PPP dynamic profile.

The remaining 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

- *Dynamic Profiles Overview*
- *Configuring Dynamic Authentication for PPP Subscribers*
- *Attaching Dynamic Profiles to Static PPP Subscriber Interfaces*
- *Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface*

classifiers (Definition)

Syntax

```
classifiers {
  type classifier-name {
    import (classifier-name | default);
    forwarding-class class-name {
      loss-priority level code-points [ aliases ] [ bit-patterns ];
    }
  }
}
```

Hierarchy Level

[edit class-of-service],
[edit class-of-service routing-instances *routing-instance-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
ieee-802.1ad option introduced in Junos OS Release 9.2.

Description

Define a CoS behavior aggregate (BA) classifier for classifying packets. You can associate the classifier with a forwarding class or code-point mapping, and import a default classifier or one that is previously defined.

NOTE: The **[edit class-of-service routing-instances *routing-instance-name*]** hierarchy level and the **dscp-ipv6** and **ieee-802.1ad** classifier types are not supported on ACX Series routers.

Options

classifier-name—Name of the aggregate behavior classifier.

type—Traffic type: **dscp**, **dscp-ipv6**, **exp**, **ieee-802.1**, **ieee-802.1ad**, **inet-precedence**.

Required Privilege Level

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

Related Documentation

- *Overview of BA Classifier Types*
- *Example: Configuring CoS for a PBB Network on MX Series Routers*
- *Configuring CoS on ACX Series Universal Access Routers*

classifiers (Logical Interface)

Syntax

```
classifiers {
   type (classifier-name | default) family (mpls | inet);
}
```

Hierarchy Level [edit class-of-service interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Apply a CoS aggregate behavior classifier to a logical interface. You can apply a default classifier or one that is previously defined.

Options ***classifier-name***—Name of the aggregate behavior classifier.

type—Traffic type.

Values: dscp, dscp-ipv6, exp, ieee-802.1, inet-precedence

NOTE: **You can only specify a family for the dscp and dscp-ipv6 types.**

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

Related Documentation

- *Default DSCP and DSCP IPv6 Classifier*
- *Applying Classifiers to Logical Interfaces*

class-of-service

Syntax class-of-service { ... }

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure Junos CoS features.

Default If you do not configure any CoS features, all packets are transmitted from output transmission queue 0.

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

Related Documentation

- *CoS Overview*

class-of-service (Dynamic Profiles)

Syntax class-of-service { ... }

Hierarchy Level [edit dynamic-profiles *profile-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Configure Junos OS CoS features in a dynamic profile.

Default If you do not configure any CoS features, all packets are transmitted from output transmission queue 0.

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

Related Documentation

- *Guidelines for Configuring Dynamic CoS for Subscriber Access*
- *Configuring Static Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access*
- *Configuring Dynamic Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access*

code-points

Syntax code-points ([*aliases*] | [*bit-patterns*]);

Hierarchy Level [edit class-of-service classifiers *type classifier-name* forwarding-class *class-name* loss-priority *level*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 8.5 for J Series devices.
Statement introduced in Junos OS Release 9.2 for SRX Series devices.

Description Specify one or more DSCP code-point aliases or bit sets for association with a forwarding class.

Options ***aliases***—Name of the DSCP alias.

bit-patterns—Value of the code-point bits, in six-bit binary form.

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

Related Documentation

- *Overview of BA Classifier Types*
- *Example: Configuring CoS for a PBB Network on MX Series Routers*
- *Example: Configuring Behavior Aggregate Classifiers*
- *Example: Configuring Forwarding Classes*

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.

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 Demux Underlying Interface*

demux-source (Dynamic Underlying Interface)

Syntax demux-source *family*;

Hierarchy Level [edit dynamic-profiles interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure the logical demultiplexing (demux) source family type on the IP demux underlying interface within a dynamic profile.

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.

dhcp-local-server

Syntax

```
dhcp-local-server {
  authentication {
    password password-string;
    username-include {
      circuit-type;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-name ;
      logical-system-name;
      mac-address;
      option-60;
      option-82 <circuit-id> <remote-id>;
      routing-instance-name;
      user-prefix user-prefix-string;
    }
  }
  dhcpv6 {
    authentication {
      ...
    }
    group group-name {
      authentication {
        ...
      }
      interface interface-name {
        exclude;
        liveness-detection {
          failure-action (clear-binding | clear-binding-if-interface-up | log-only);
          method {
            bfd {
              version (0 | 1 | automatic);
              minimum-interval milliseconds;
              minimum-receive-interval milliseconds;
              multiplier number;
              no-adaptation;
              transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
              }
              detection-time {
                threshold milliseconds;
              }
              session-mode(automatic | multihop | singlehop);
              holddown-interval milliseconds;
            }
          }
        }
        overrides {
          interface-client-limit number;
          process-inform {
            pool pool-name;
          }
```

```
            rapid-commit;
        }
        service-profile dynamic-profile-name;
        trace;
        upto upto-interface-name;
    }
    overrides {
        delegated-pool;
        interface-client-limit number;
        process-inform {
            pool pool-name;
        }
        rapid-commit;
    }
    route-suppression;
    service-profile dynamic-profile-name;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            detection-time {
                threshold milliseconds;
            }
            session-mode(automatic | multihop | singlehop);
            holddown-interval milliseconds;
        }
    }
}
overrides {
    delegated-pool;
    interface-client-limit number;
    process-inform {
        pool pool-name;
    }
    rapid-commit;
}
reconfigure {
    attempts attempt-count;
    clear-on-abort;
    strict;
    timeout timeout-value;
    token token-value;
    trigger {
        radius-disconnect;
    }
}
```

```
    route-suppression;
    service-profile dynamic-profile-name;
  }
  duplicate-clients-on-interface;
  dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary
    primary-profile-name>;
  forward-snooped-clients (all-interfaces | configured-interfaces |
    non-configured-interfaces);
  group group-name {
    authentication {
      ...
    }
    dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary
      primary-profile-name>;
    interface interface-name {
      exclude;
      liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
          bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
              minimum-interval milliseconds;
              threshold milliseconds;
            }
            detection-time {
              threshold milliseconds;
            }
            session-mode(automatic | multihop | singlehop);
            holddown-interval milliseconds;
          }
        }
      }
      overrides {
        client-discover-match <option60-and-option82>;
        interface-client-limit number;
        process-inform {
          pool pool-name;
        }
      }
      service-profile dynamic-profile-name;
      trace;
      upto upto-interface-name;
    }
    overrides {
      client-discover-match <option60-and-option82>;
      interface-client-limit number;
      process-inform {
        pool pool-name;
      }
    }
    route-suppression;
```

```
    service-profile dynamic-profile-name;
  }
  liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
      bfd {
        version (0 | 1 | automatic);
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
          minimum-interval milliseconds;
          threshold milliseconds;
        }
        detection-time {
          threshold milliseconds;
        }
        session-mode(automatic | multihop | singlehop);
        holddown-interval milliseconds;
      }
    }
  }
  overrides {
    client-discover-match <option60-and-option82>;
    interface-client-limit number;
    process-inform {
      pool pool-name;
    }
  }
  pool-match-order {
    external-authority;
    ip-address-first;
    option-82;
  }
  reconfigure {
    attempts attempt-count;
    clear-on-abort;
    strict;
    timeout timeout-value;
    token token-value;
    trigger {
      radius-disconnect;
    }
  }
  route-suppression;
  service-profile dynamic-profile-name;
}
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description

Configure Dynamic Host Configuration Protocol (DHCP) local server options on the router or switch and enable the router or switch to function as an extended DHCP local server. The DHCP local server receives DHCP request and reply packets from DHCP clients and then responds with an IP address and other optional configuration information to the client.

The extended DHCP local server is incompatible with the DHCP server on J Series routers and so is not supported on J Series routers. Also, the DHCP local server and the DHCP/BOOTP relay server, which are configured under the **[edit forwarding-options helpers]** hierarchy level, cannot both be enabled on the router or switch at the same time. The extended DHCP local server is fully compatible with the extended DHCP relay feature.

The **dhcpv6** stanza configures the router or switch to support Dynamic Host Configuration Protocol for IPv6 (DHCPv6). The DHCPv6 local server is fully compatible with the extended DHCP local server and the extended DHCP relay feature.

NOTE: **When you configure the dhcp-local-server statement at the routing instance hierarchy level, you must use a routing instance type of virtual-router.**

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Extended DHCP Local Server Overview*
- *DHCPv6 Local Server Overview*
- *Configuring a DHCP Server on EX Series Switches (CLI Procedure)*

domain-name (DHCP Local Server)

Syntax

domain-name *domain-name-string*;

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server authentication username-include],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 authentication username-include],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 group *group-name* authentication username-include],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* authentication username-include],
[edit logical-systems *logical-system-name* system services dhcp-local-server authentication username-include],
[edit logical-systems *logical-system-name* system services dhcp-local-server dhcpv6 authentication username-include],
[edit logical-systems *logical-system-name* system services dhcp-local-server dhcpv6 group *group-name* authentication username-include],
[edit logical-systems *logical-system-name* system services dhcp-local-server group *group-name* authentication username-include],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server authentication username-include],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 authentication username-include],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 group *group-name* authentication username-include],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* authentication username-include],
[edit routing-instances *routing-instance-name* system services dhcp-local-server authentication username-include],
[edit routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 authentication username-include],
[edit routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 group *group-name* authentication username-include],
[edit routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* authentication username-include],
[edit system services dhcp],
[edit system services dhcp-local-server authentication username-include],
[edit system services dhcp-local-server dhcpv6 authentication username-include],
[edit system services dhcp-local-server dhcpv6 group *group-name* authentication username-include],
[edit system services dhcp-local-server group *group-name* authentication username-include]

Release Information

Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Specify the domain name that is concatenated with the username during the subscriber authentication or DHCP client authentication process.

Options

domain-name-string—Domain name formatted string.

Required Privilege Level

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

Related Documentation

- *Using External AAA Authentication Services with DHCP*

dynamic-profile (Stacked VLAN)

Syntax

```
dynamic-profile profile-name {
  accept (any | dhcp-v4 |dhcp-v6| inet | inet6 | pppoe);
  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.

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*
- *Associating a Stacked VLAN Dynamic Profile with an Interface*

dynamic-profile (DHCP Local Server)

Syntax

```
dynamic-profile profile-name {
   aggregate-clients (merge | replace);
   use-primary primary-profile-name;
}
```

Hierarchy Level

[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server dhcpv6 group *group-name*],
[edit system services dhcp-local-server dhcpv6 group *group-name* interface *interface-name*],
[edit system services dhcp-local-server group *group-name*],
[edit system services dhcp-local-server group *group-name* interface *interface-name*],
[edit logical-systems *logical-system-name* system services dhcp-local-server ...],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server ...],
[edit routing-instances *routing-instance-name* system services dhcp-local-server ...]

Release Information

Statement introduced in Junos OS Release 9.2.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Options **aggregate-clients** and **use-primary** introduced in Junos OS Release 9.3.
Support at the **[edit ... interface]** hierarchy levels introduced in Junos OS Release 11.2.

Description

Specify the dynamic profile that is attached to all interfaces, a named group of interfaces, or a specific interface.

Options

profile-name—Name of the dynamic profile.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces*
- *Configuring a Default Subscriber Service*

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);
          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;
          match-order [match-order];
```

```
    term term-name {
      from {
        match-conditions;
      }
      then {
        action;
        action-modifiers;
      }
      only-at-create;
    }
  }
  firewall {
    family family {
      fast-update-filter filter-name {
        interface-specific;
      match-order [match-order];
        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 policer-name {
      aggregate {
        if-exceeding {
          bandwidth-limit-limit bps;
          burst-size-limit bytes;
        }
        then {
```

```
                policer-action;
              }
            }
            premium {
              if-exceeding {
                bandwidth-limit bps;
                burst-size-limit bytes;
              }
              then {
                policer-action;
              }
            }
          }
          three-color-policer policer-name {
            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;
              }
            }
          }
        }
        policy-options {
          prefix-listname {
            ip-addresses;
          }
        }
      }
    }
    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 {
        auto-configure {
          agent-circuit-identifier {
            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;
      }
      output filter-name {
        precedence precedence;
      }
    }
    rpf-check {
      fail-filter filter-name;
      mode loose;
    }
    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;
```

```
            }
          }
          unnumbered-address interface-name <preferred-source-address address>;
        }
        ppp-options {
          chap;
          pap;
        }
        vlan-id number;
        vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
      }
    }
    interfaces {
      demux0 {...}
    }
    interfaces {
      pp0 {...}
    }
    protocols {
      igmp {
        interface interface-name {
          accounting;
          disable;
          group-policy;
          immediate-leave
          no-accounting;
          promiscuous-mode;
          ssm-map ssm-map-name;
          static {
            group group {
              source source;
            }
          }
          version version;
        }
        mld {
          interfaceinterface-name {
            disable;
            (accounting | no-accounting);
            group-policy;
            immediate-leave;
            oif-map;
            passive;
            ssm-map ssm-map-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;
            }
        }
        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;
            }
        }
        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;
            }
          }
          access-internal {
            route subscriber-ip-address {
              qualified-next-hop underlying-interface {
                mac-address address;
              }
            }
          }
          multicast {
            interface interface-name {
              no-qos-adjust;
            }
          }
        }
        variables {
          variable-name {
            default-value default-value;
            equals expression;
            mandatory;
            radius {
              vendor-id id {
                attribute attribute-number;
                tag tag-number;
              }
            }
            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.

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*

family

Syntax

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

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

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

Options *family*—Protocol family:

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

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

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

The remaining statements are explained separately.

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

Related Documentation

- *Configuring the Protocol Family*
- *Example: Configuring E-LINE and E-LAN Services for a PBB Network on MX Series Routers*
- *Junos OS Services Interfaces Library for Routing Devices*

family (Address-Assignment Pools)

Syntax

```
family family {
  dhcp-attributes {
    [protocol-specific attributes]
  }
  host hostname {
    hardware-address mac-address;
    ip-address ip-address;
  }
  network ip-prefix/<prefix-length>;
  prefix ipv6-prefix;
  range range-name {
    high upper-limit;
    low lower-limit;
    prefix-length prefix-length;
  }
}
```

Hierarchy Level

[edit access address-assignment pool *pool-name*]

Release Information

Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.3 for EX Series switches.

Description

Configure the protocol family for the address-assignment pool.

NOTE: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Options

family—Protocol family:

- **inet**—Internet Protocol version 4 suite
- **inet6**—Internet Protocol version 6 suite

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Address-Assignment Pools Overview*
- *Configuring Address-Assignment Pools*

family (Dynamic Standard Interface)

Syntax

```
family family {
  access-concentrator name;
  address address;
  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;
    }
    output filter-name {
      precedence precedence;
    }
  }
  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>;
  unnumbered-address interface-name <preferred-source-address address>;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information

Statement introduced in Junos OS Release 9.2.
Option **pppoe** introduced 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.

Required Privilege Level

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

Related Documentation

- For general information about configuring static interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*.
- "Configuring the Protocol Family," in the *Junos OS Network Interfaces Library for Routing Devices*.

firewall

Syntax

```
firewall {
   atm-policer atm-policer-name {
      ... atm-policer-configuration ...
   }
   family protocol-family-name {
      ... protocol-family-configuration ...
   }
   filter ipv4-filter-name {
      ... ipv4-filter-configuration ...
   }
   hierarchical-policer hierarchical-policer-name {
      ... hierarchical-policer-configuration ...
   }
   interface-set interface-set-name {
      ... interface-set-configuration ...
   }
   policer two-color-policer-name {
      ... two-color-policer-configuration ...
   }
   three-color-policer three-color-policer-name {
      ... three-color-policer-configuration ...
   }
}
```

Hierarchy Level

[edit],
[edit logical-systems *logical-system-name*]
[edit dynamic-profiles *profile-name*],

Release Information

Statement introduced before Junos OS Release 7.4.
Logical systems support introduced in Junos OS Release 9.3.

Description

Configure firewall filters.

The statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Guidelines for Configuring Firewall Filters*
- *Guidelines for Configuring Service Filters*
- *Guidelines for Configuring Simple Filters*
- *Configuring Multifield Classifiers*
- *Using Multifield Classifiers to Set PLP*

filter (Configuring)

Syntax

```
filter filter-name {
   accounting-profile name;
   enhanced-mode;
   interface-shared;
   interface-specific;
   physical-interface-filter;
   term term-name {
      ... term configuration ...
   }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* firewall family *family-name*],
[edit firewall family *family-name*],
[edit logical-systems *logical-system-name* firewall family *family-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Logical systems support introduced in Junos OS Release 9.3.
physical-interface-filter statement introduced in Junos OS Release 9.6.
Support at the **[edit dynamic-profiles ... family family-name]** hierarchy level introduced in Junos OS Release 11.4.
Support for the **interface-shared**> statement introduced in Junos OS Release 12.2.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure firewall filters.

Options

filter-name—Name that identifies the filter. This must be a non-reserved string of not more than 64 characters. To include spaces in the name, enclose it in quotation marks (“ ”). In Junos OS Release 9.0 and later, you can no longer use special characters within the name of a firewall filter. Firewall filter names are restricted from having the form **__.*__** (beginning and ending with underscores) or **__.*** (beginning with an underscore.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Guidelines for Configuring Firewall Filters*
- *Guidelines for Applying Firewall Filters*
- *Configuring Multifield Classifiers*
- *Using Multifield Classifiers to Set PLP*
- *simple-filter (Configuring)*

filter (Dynamic Interface Unit)

Syntax

```
filter {
   input filter-name;
   output filter-name;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number*],
[edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number*],

Release Information

Statement introduced in Junos OS Release 9.2.

Description

Apply a dynamic filter to an interface, regardless of its family type.

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.

The remaining 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

- For general information about configuring firewall filters, see the *Routing Policy Feature Guide for Routing Devices*
- *Dynamic Firewall Filters Overview*
- *Classic Filters Overview*
- *Basic Classic Filter Syntax*
- *Dynamically Attaching Statically Created Filters for Any Interface Type*

forwarding-class (BA Classifiers)

Syntax

```
forwarding-class class-name {
  loss-priority level code-points [ aliases ] [ bit-patterns ];
}
```

Hierarchy Level [edit class-of-service classifiers *type classifier-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define forwarding class name and option values.

Options ***class-name***—Name of the forwarding class.

The remaining statements are explained separately.

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

Related Documentation

- *Defining Classifiers*
- *Example: Configuring CoS for a PBB Network on MX Series Routers*

forwarding-class (Interfaces)

Syntax forwarding-class *class-name*;

Hierarchy Level [edit class-of-service interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for ACX Series routers.

Description Associate a forwarding class configuration or default mapping with a specific interface.

Options ***class-name***—Name of the forwarding class.

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

Related Documentation

- *Applying Forwarding Classes to Interfaces*

forwarding-classes (Class-of-Service)

Syntax

```
forwarding-classes {
   class queue-num queue-number priority (high | low);
   queue queue-number class-name priority (high | low) [ policing-priority (premium | normal) ];
}
```

Hierarchy Level

[edit class-of-service]

Release Information

Statement introduced before Junos OS Release 7.4.
policing-priority option introduced in Junos OS Release 9.5.
Statement introduced on PTX Series Packet Transport Routers in Junos OS Release 12.1.

Description

Associate the forwarding class with a queue name and number. For M320, MX Series, T Series routers and EX Series switches only, you can configure fabric priority queuing by including the **priority** statement. For Enhanced IQ PICs, you can include the **policing-priority** option.

NOTE: **The priority add policing-priority options are not supported on PTX Series Packet Transport Routers.**

The statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Configuring Forwarding Classes*
- *Forwarding Classes and Fabric Priority Queues*
- *Example: Configuring CoS for a PBB Network on MX Series Routers*
- *Configuring Layer 2 Policers on IQE PICs*
- *Classifying Packets by Egress Interface*

grace-period

Syntax grace-period *seconds*;

Hierarchy Level [edit access address-assignment pool *pool-name* family (inet | inet6) dhcp-attributes]

Release Information Statement introduced in Junos OS Release 9.0.

Description Configure the amount of time that the client retains the address lease after the lease expires. The address cannot be reassigned to another client during the grace period.

Options ***seconds***—Number of seconds the lease is retained.

Range: 0 through 4,294,967,295 seconds

Default: 0 (no grace period)

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

Related Documentation

- *Configuring Address-Assignment Pools*

group (DHCP Local Server)

Syntax

```
group group-name {
  authentication {
    password password-string;
    username-include {
      circuit-type;
      client-id;
      delimiter delimiter-character;
      domain-name domain-name-string;
      logical-system-name;
      mac-address;
      option-60;
      option-82 <circuit-id> <remote-id>;
      relay-agent-interface-id
      relay-agent-remote-id;
      relay-agent-subscriber-id;
      routing-instance-name;
      user-prefix user-prefix-string;
    }
  }
  dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary
    primary-profile-name>;
  interface interface-name {
    exclude;
    overrides {
      client-discover-match <option60-and-option82>;
      interface-client-limit number;
      process-inform {
        pool pool-name;
      }
      rapid-commit;
    }
    service-profile dynamic-profile-name;
    trace;
    upto upto-interface-name;
  }
  liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
      bfd {
        version (0 | 1 | automatic);
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
          minimum-interval milliseconds;
          threshold milliseconds;
        }
        detection-time {
          threshold milliseconds;
        }
        session-mode(automatic | multihop | singlehop);
```

```
            holddown-interval milliseconds;
          }
        }
      }
      overrides {
        client-discover-match <option60-and-option82>;
        delegated-pool;
        interface-client-limit number;
        process-inform {
          pool pool-name;
        }
        rapid-commit;
      }
      reconfigure {
        attempts attempt-count;
        clear-on-abort;
        strict;
        timeout timeout-value;
        token token-value;
        trigger {
          radius-disconnect;
        }
      }
      route-suppression;
      service-profile dynamic-profile-name;
    }
```

Hierarchy Level

[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server ...],
[edit logical-systems *logical-system-name* system services dhcp-local-server ...],
[edit routing-instances *routing-instance-name* system services dhcp-local-server ...]

Release Information

Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description

Configure a group of interfaces that have a common configuration, such as authentication parameters. A group must contain at least one interface.

Options

group-name—Name of the group.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Extended DHCP Local Server Overview*
- *Grouping Interfaces with Common DHCP Configurations*
- *Using External AAA Authentication Services with DHCP*
- *Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces*
- *Configuring a DHCP Server on EX Series Switches (CLI Procedure)*

immediate-update

Syntax

immediate-update;

Hierarchy Level

[edit access profile *profile-name* accounting]

Release Information

Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.

Description

Configure the router or switch to send an Acct-Update message to the RADIUS accounting server on receipt of a response (for example, an ACK or timeout) to the Acct-Start message.

Required Privilege Level

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

Related Documentation

- *Configuring RADIUS Server Parameters for Subscriber Access*
- *Configuring Per-Subscriber Session Accounting*

interface (DHCP Local Server)

Syntax

```
interface interface-name {
  exclude;
  overrides {
    client-discover-match <option60-and-option82>;
    interface-client-limit number;
    rapid-commit;
  }
  service-profile dynamic-profile-name;
  trace;
  upto upto-interface-name;
}
```

Hierarchy Level

[edit system services dhcp-local-server group *group-name*],
[edit system services dhcp-local-server dhcpv6 group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server ...],
[edit logical-systems *logical-system-name* system services dhcp-local-server ...],
[edit routing-instances *routing-instance-name* system services dhcp-local-server ...]

Release Information

Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Options **upto** and **exclude** introduced in Junos OS Release 9.1.

Description

Specify one or more interfaces, or a range of interfaces, that are within a specified group on which the DHCP local server is enabled. You can repeat the **interface** ***interface-name*** statement to specify multiple interfaces within a group, but you cannot specify the same interface in more than one group. Also, you cannot use an interface that is being used by the DHCP relay agent.

NOTE: DHCP values are supported in Integrated Routing and Bridging (IRB) configurations. When you configure an IRB interface in a network that is using DHCP, the DHCP information (for example, authentication, address assignment, and so on) is propagated in the associated bridge domain. This enables the DHCP server to configure client IP addresses residing within the bridge domain. IRB currently only supports static DHCP configurations. For additional information about how to configure IRB, see the *Ethernet Networking Feature Guide for MX Series Routers*.

Options

exclude—Exclude an interface or a range of interfaces from the group. This option and the **overrides** option are mutually exclusive.

interface-name—Name of the interface. You can repeat this option multiple times.

upto-interface-name—Upper end of the range of interfaces; the lower end of the range is the interface-name entry. The interface device name of the ***upto-interface-name*** must be the same as the device name of the ***interface-name***.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Extended DHCP Local Server Overview*
- *Grouping Interfaces with Common DHCP Configurations*
- *Using External AAA Authentication Services with DHCP*

interfaces

Syntax

interfaces { ... }

Hierarchy Level

[edit]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure interfaces on the router or switch.

Default

The management and internal Ethernet interfaces are automatically configured. You must configure all other interfaces.

Required Privilege Level

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

Related Documentation

- *Physical Interface Configuration Statements Overview*
- *Configuring Aggregated Ethernet Link Protection*

interfaces

Syntax

```
interfaces {
  interface-name {
    classifiers{
      dscp(classifier-name | default) {
    }
      ieee-802.1 (classifier-name | default)  vlan-tag (inner | outer | classifier-name);
      inet-precedence (rewrite-name | default);
    }
    input-scheduler-map map-name;
    input-shaping-rate rate;
    irb {
      unit logical-unit-number {
        classifiers {
          type (classifier-name | default);
        }
        rewrite-rules {
          dscp (rewrite-name | default);
          dscp-ipv6 (rewrite-name | default);
          exp (rewrite-name | default)protocol protocol-types;
          ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
          inet-precedence (rewrite-name | default);
        }
      }
    }
    member-link-scheduler (replicate | scale);
    rewrite-rules {
      dscp (rewrite-name | default);
      ieee-802.1 (rewrite-name | default) vlan-tag (outer);
      inet-precedence (rewrite-name | default);
      }
    }
    scheduler-map map-name;
    scheduler-map-chassis map-name;
    shaping-rate rate;
    unit logical-unit-number {
      classifiers {
        type (classifier-name | default) family (mpls | inet);
      }
      forwarding-class class-name;
      fragmentation-map map-name;
      input-shaping-rate (percent percentage | rate);
      input-traffic-control-profile profile-name shared-instance instance-name;
      output-traffic-control-profile profile-name shared-instance instance-name;
      per-session-scheduler;
      rewrite-rules {
        dscp (rewrite-name | default);
        dscp-ipv6 (rewrite-name | default);
        exp (rewrite-name | default) protocol protocol-types;
        exp-push-push-push default;
        exp-swap-push-push default;
        ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
        inet-precedence (rewrite-name | default);
```

```
            }
            scheduler-map map-name;
            shaping-rate rate;
            translation-table (to-dscp-from-dscp | to-dscp-ipv6-from-dscp-ipv6 | to-exp-from-exp
              | to-inet-precedence-from-inet-precedence) table-name;
          }
        }
        interface-set interface-set-name {
          excess-bandwidth-share;
          internal-node;
          output-traffic-control-profile profile-name;
          output-traffic-control-profile-remaining profile-name;
        }
      }
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced before Junos OS Release 7.4.
Interface-set level added in Junos OS Release 8.5.

Description Configure interface-specific CoS properties for incoming packets.

NOTE: The **dscp-ipv6** and **ieee-802.1ad** classifier types are not supported on ACX Series routers. For further information about support on ACX Series routers, see *Understanding CoS CLI Configuration Statements on ACX Series Universal Access Routers*.

Options The remaining statements are explained separately.

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

Related Documentation

- *Overview of BA Classifier Types*
- *Configuring Rewrite Rules*
- *Understanding CoS CLI Configuration Statements on ACX Series Universal Access Routers*

interfaces (Static and Dynamic Subscribers)

Syntax

```
interfaces {
  interface-name {
    unit logical-unit-number {
      auto-configure {
        agent-circuit-identifier {
          dynamic-profile profile-name;
        }
      }
      family family {
        access-concentrator name;
        address address;
        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;
          shared-name filter-shared-name;
          output filter-name;
          shared-name filter-shared-name;
        }
        ppp-options {
          chap;
          pap;
        }
        proxy-arp;
        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 {
          access-concentrator name;
          address address;
          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 {
          authentication [ authentication-protocols ];
          chap {
            challenge-length minimum minimum-length maximum maximum-length;
          }
          pap;
        }
        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.

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.

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

Related Documentation

- *Configuring Static Subscriber Interfaces in Dynamic Profiles*
- *Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles*
- *Configuring Dynamic PPPoE Subscriber Interfaces Using Dynamic Profiles*
- *Configuring Dynamic VLANs Based on Agent Circuit Identifier Information*
- *Subscriber Interface Overview*
- *Relationship Between Subscribers and Interfaces in an Access Network*
- For general information about configuring static interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*
- For information about static IP demux interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*

interfaces (Dynamic CoS Definition)

Syntax

```
interfaces {
  interface-name {
    unit logical-unit-number {
      classifiers {
        dscp (classifier-name | default);
        dscp-ipv6 (classifier-name | default);
        ieee-802.1 (classifier-name | default) vlan-tag (inner | outer)
        inet-precedence (classifier-name | default);
      }
      output-traffic-control-profile (profile-name | $junos-cos-traffic-control-profile);
      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);
      }
    }
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* class-of-service]

Release Information

Statement introduced in Junos OS Release 9.2.

Description

Configure interface-specific CoS properties for incoming packets.

Options

interface-name—Either the specific name of the interface you want to assign to the dynamic profile or the interface variable ($junos-interface-ifd-name). The interface variable is dynamically replaced with the interface the client accesses when connecting to the router.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Guidelines for Configuring Dynamic CoS for Subscriber Access*
- *Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile*

interface-specific (Firewall Filters)

Syntax interface-specific;

Hierarchy Level [edit firewall family *family-name* filter *filter-name*],
[edit logical-systems *logical-system-name* firewall family *family-name* filter *filter-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Logical systems support introduced in Junos OS Release 9.3.

Description Configure interface-specific names for firewall counters.

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

Related Documentation
- *Configuring Firewall Filters and Policers for VPLS*
- *Interface-Specific Firewall Filter Instances Overview*

loss-priority (BA Classifiers)

Syntax loss-priority *level*;

Hierarchy Level [edit class-of-service classifiers *type classifier-name* forwarding-class *class-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify packet loss priority value for a specific set of code-point aliases and bit patterns.

Options ***level*** can be one of the following:

- **high**—Packet has high loss priority.
- **medium-high**—Packet has medium-high loss priority.
- **medium-low**—Packet has medium-low 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
- *Overview of BA Classifier Types*
- *Example: Configuring CoS for a PBB Network on MX Series Routers*
- *Configuring Tricolor Marking*

maximum-lease-time

Syntax maximum-lease-time *seconds*;

Hierarchy Level [edit access address-assignment pool *pool-name* family (inet | inet6) dhcp-attributes]

Release Information Statement introduced in Junos OS Release 9.0.

Description Specify the maximum length of time, in seconds, that the lease is held for a client if the client does not renew the lease. This is equivalent to DHCP option 51.

Options ***seconds***—Maximum number of seconds the lease can be held.

Range: 30 through 4,294,967,295 seconds

Default: 86,400 (24 hours)

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

Related Documentation

- *Configuring Address-Assignment Pools*

network

Syntax network *ip-prefix*</*prefix-length*>;

Hierarchy Level [edit access address-assignment pool *pool-name* family inet]

Release Information Statement introduced in Junos OS Release 9.0.

Description Configure subnet information for an IPv4 address-assignment pool.

Options ***ip-prefix***—IP version 4 address or prefix value.

prefix-length—(Optional) Subnet mask.

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

Related Documentation

- *Configuring Address-Assignment Pools*

order

Syntax order [*accounting-method*];

Hierarchy Level [edit access profile *profile-name* accounting]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.

Description Set the order in which the Junos OS tries different accounting methods for client activity. When a client logs in, the software tries the accounting methods in the specified order.

Options ***accounting-method***—One or more accounting methods. When a client logs in, the software tries the accounting methods in the following order, from first to last. The only valid value is **radius** for RADIUS accounting.

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

Related Documentation
- *Configuring Authentication and Accounting Parameters for Subscriber Access*

output-traffic-control-profile (Dynamic CoS Definition)

Syntax output-traffic-control-profile (*profile-name* | $junos-cos-traffic-control-profile);

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.2.
Variable **$junos-cos-traffic-control-profile** introduced in Junos OS Release 11.2.

Description Apply an output traffic scheduling and shaping profile to the logical interface.

Options ***profile-name***—Name of the traffic-control profile to be applied to this interface

$junos-cos-traffic-control-profile—Variable for the traffic-control profile that is specified for the logical interface. The variable is replaced with the traffic-control profile when the subscriber is authenticated at login.

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

Related Documentation
- *Guidelines for Configuring Dynamic CoS for Subscriber Access*
- *Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile*
- *Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers*
- traffic-control-profiles on page 151

pap (Dynamic PPP)

Syntax pap;

Hierarchy Level [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 in Junos OS Release 9.5.
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 12.2.

Description Specify PAP authentication in a PPP dynamic profile.

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 Dynamic Authentication for PPP Subscribers*
- *Attaching Dynamic Profiles to Static PPP Subscriber Interfaces*
- *Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface*

password (DHCP Local Server)

Syntax

password *password-string*;

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server authentication],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 authentication],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 group *group-name* authentication],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* authentication],
[edit logical-systems *logical-system-name* system services dhcp-local-server authentication],
[edit logical-systems *logical-system-name* system services dhcp-local-server dhcpv6 authentication],
[edit logical-systems *logical-system-name* system services dhcp-local-server dhcpv6 group *group-name* authentication],
[edit logical-systems *logical-system-name* system services dhcp-local-server group *group-name* authentication],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server authentication],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 authentication],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 group *group-name* authentication],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* authentication],
[edit routing-instances *routing-instance-name* system services dhcp-local-server authentication],
[edit routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 authentication],
[edit routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 group *group-name* authentication],
[edit routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* authentication],
[edit system services dhcp-local-server authentication],
[edit system services dhcp-local-server dhcpv6],
[edit system services dhcp-local-server dhcpv6 group *group-name* authentication],
[edit system services dhcp-local-server group *group-name* authentication]

Release Information

Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure the password that is sent to the external AAA authentication server for subscriber authentication or DHCP client authentication.

Options

password-string—Authentication password.

Required Privilege Level

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

Related Documentation

- *Using External AAA Authentication Services with DHCP*

pool (Address-Assignment Pools)

Syntax

```
pool pool-name {
  family family {
    dhcp-attributes {
      [ protocol-specific attributes ]
    }
    host hostname {
      hardware-address mac-address;
      ip-address ip-address;
    }
    network ip-prefix/<prefix-length>;
    prefix ipv6-prefix;
    range range-name {
      high upper-limit;
      low lower-limit;
      prefix-length prefix-length;
    }
  }
  link pool-name;
}
```

Hierarchy Level

[edit access address-assignment]

Release Information

Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description

Configure the name of an address-assignment pool.

NOTE: Subordinate statement support depends on the platform. See individual statement topics for more detailed support information.

Options

pool-name—Name assigned to the address-assignment pool.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Address-Assignment Pools Overview*
- *Configuring Address-Assignment Pools*
- *Configuring a DHCP Server on EX Series Switches (CLI Procedure)*

pool-match-order

Syntax

```
pool-match-order {
   external-authority;
   ip-address-first;
   option-82;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server],
[edit logical-systems *logical-system-name* system services dhcp-local-server],
[edit routing-instances *routing-instance-name* system services dhcp-local-server],
[edit system services dhcp-local-server]

Release Information

Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.1.

Description

Configure the order in which the DHCP local server uses information in the DHCP client PDU to determine how to obtain an address for the client.

The remaining statements are explained separately.

Default

DHCP local server uses the **ip-address-first** method to determine which address pool to use.

Required Privilege Level

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

Related Documentation

- *Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use*
- *Extended DHCP Local Server Overview*
- *Configuring a DHCP Server on EX Series Switches (CLI Procedure)*

ppp-options (Dynamic PPP)

Syntax

```
ppp-options {
  authentication [ authentication-protocols ];
  chap {
    challenge-length minimum minimum-length maximum maximum-length;
  }
  on-demand-ip-address;
  pap;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces pp0 unit “$junos-interface-unit”],
[edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit “$junos-interface-unit”]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles *profile-name* interfaces "$junos-interface-ifd-name" unit “$junos-interface-unit”]** hierarchy level introduced in Junos OS Release 12.2.

Description

Configure PPP-specific interface properties in a dynamic profile.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Dynamic Profiles Overview*
- *Configuring Dynamic Authentication for PPP Subscribers*
- *Attaching Dynamic Profiles to Static PPP Subscriber Interfaces*
- *Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface*

pppoe-options (Dynamic PPPoE)

Syntax

```
pppoe-options {
   underlying-interface interface-name;
   server;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces pp0 unit “$junos-interface-unit”]

Release Information

Statement introduced in Junos OS Release 10.1.

Description

Configure the underlying interface and PPPoE server mode for a dynamic PPPoE logical interface in a dynamic profile.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Configuring a Basic PPPoE Dynamic Profile*
- For information about creating static PPPoE interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*

priority (Fabric Queues, Schedulers)

Syntax priority (high | low)scheduler *scheduler-name*;

Hierarchy Level [edit class-of-service fabric scheduler-map]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced before Junos OS 11.4 for EX Series switches.

Description Define Fabric traffic priority. For M320, MX Series, T Series routers and EX Series switches only, specify the fabric priority with which a scheduler is associated.

For a scheduler that you associate with a fabric priority, you cannot include the **buffer-size**, **transmit-rate**, or **priority** statements at the **[edit class-of-service schedulers *scheduler-name*]** hierarchy level.

On EX Series switches, this statement is supported only on EX8200 standalone switches and EX8200 Virtual Chassis.

Options **high**—Scheduler has high priority.

low—Scheduler has low priority.

The remaining statements are explained separately.

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

Related Documentation

- See *Associating Schedulers with Fabric Priorities*.
- *Understanding Junos OS CoS Components for EX Series Switches*

profile (Access)

Syntax

```
profile profile-name {
  accounting {
    address-change-immediate-update
    accounting-stop-on-access-deny;
    accounting-stop-on-failure;
    coa-immediate-update;
    coa-no-override service-class-attribute;
    duplication;
    duplication-vrf {
      access-profile-name profile-name;
      vrf-name vrf-name;
    }
    immediate-update;
    order [ accounting-method ];
    send-acct-status-on-config-change;
    statistics (time | volume-time);
    update-interval minutes;
    wait-for-acct-on-ack;
    }
  authentication-order [ authentication-methods ];
  client client-name {
    chap-secret chap-secret;
    group-profile profile-name;
    ike {
      allowed-proxy-pair {
        remote remote-proxy-address local local-proxy-address;
      }
      pre-shared-key (ascii-text character-string | hexadecimal hexadecimal-digits);
      ike-policy policy-name;
      interface-id string-value;
    }
    l2tp {
      aaa-access-profile profile-name;
      interface-id interface-id;
      lcp-renegotiation;
      local-chap;
      maximum-sessions-per-tunnel number;
      multilink {
        drop-timeout milliseconds;
        fragment-threshold bytes;
      }
      ppp-authentication (chap | pap);
      ppp-profile profile-name;
      shared-secret shared-secret;
    }
    pap-password pap-password;
    ppp {
      cell-overhead;
      encapsulation-overhead bytes;
      framed-ip-address ip-address;
      framed-pool framed-pool;
      idle-timeout seconds;
```

```
        interface-id interface-id;
        keepalive seconds;
        primary-dns primary-dns;
        primary-wins primary-wins;
        secondary-dns secondary-dns;
        secondary-wins secondary-wins;
      }
      user-group-profile profile-name;
    }
    domain-name-server;
    domain-name-server-inet;
    domain-name-server-inet6;
    provisioning-order (gx-plus | jsrc);
    radius {
      accounting-server [ ip-address ];
      authentication-server [ ip-address ];
      options {
        accounting-session-id-format (decimal | description);
        calling-station-id-delimiter delimiter-character;
        calling-station-id-format {
          agent-circuit-id;
          agent-remote-id;
          interface-description;
          nas-identifier;
        }
        client-accounting-algorithm (direct | round-robin);
        client-authentication-algorithm (direct | round-robin);
        coa-dynamic-variable-validation;
        ethernet-port-type-virtual;
        interface-description-format {
          exclude-adapter;
          exclude-sub-interface;
        }
        juniper-dsl-attributes;
        nas-identifier identifier-value;
        nas-port-extended-format {
          adapter-width width;
          ae-width width;
          port-width width;
          slot-width width;
          stacked-vlan-width width;
          vlan-width width;
          atm {
            adapter-width width;
            port-width width:
            slot-width width;
            vci-width width:
            vpi-width width;
          }
        }
        nas-port-id-delimiter delimiter-character;
        nas-port-id-format {
          agent-circuit-id;
          agent-remote-id;
          interface-description;
          nas-identifier;
```

```
            }
            nas-port-type {
              ethernet {
                port-type;
              }
            }
            revert-interval interval;
            vlan-nas-port-stacked-format;
          }
          attributes {
            exclude {
              ...
            }
            ignore {
              framed-ip-netmask;
              input-filter;
              logical-system:routing-instance;
              output-filter;
            }
          }
        }
        radius-server server-address {
          accounting-port port-number;
          port port-number;
          retry attempts;
          routing-instance routing-instance-name;
          secret password;
          max-outstanding-requests value;
          source-address source-address;
          timeout seconds;
        }
        service {
          accounting-order (activation-protocol | radius);
        }
        session-options {
          client-group [ group-names ];
          client-idle-timeout minutes;
          client-session-timeoutminutes;
        }
      }
```

Hierarchy Level [edit access]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure PPP CHAP, or a profile and its subscriber access, L2TP, or PPP properties.

Options *profile-name*—Name of the profile.

For CHAP, the name serves as the mapping between peer identifiers and CHAP secret keys. This entity is queried for the secret key whenever a CHAP challenge or response is received.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Configuring the PPP Authentication Protocol*
- *Configuring Access Profiles for L2TP or PPP Parameters*
- *Configuring L2TP Properties for a Client-Specific Profile*
- *Configuring an L2TP LNS with Inline Service Interfaces*
- *Configuring PPP Properties for a Client-Specific Profile*
- *Configuring Service Accounting with JSRC*
- *AAA Service Framework Overview*
- show network-access aaa statistics on page 178
- *clear network-access aaa statistics*

processes

Syntax

```
processes {
  process-name (enable | disable) failover (alternate-media | other-routing-engine);
  timeout seconds;
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure which Junos OS processes are running on the router or switch.

CAUTION: **Never disable any of the software processes unless instructed to do so by a customer support engineer.**

Default All processes are enabled by default.

Options

(enable | disable)—(Optional) Enable or disable a specified process.

failover (alternate-media | other-routing-engine)—(Optional) For routers or switches with redundant Routing Engines only, switch to backup media if a process fails repeatedly. If a process fails four times within 30 seconds, the router or switch reboots from the alternate media or the other Routing Engine.

process-name—One of the valid process names. You can obtain a complete list of process names by using the CLI command completion feature. After specifying a process name, command completion also indicates any additional options for that process.

timeout *seconds*—(Optional) How often the system checks the watchdog timer, in seconds. If the watchdog timer has not been checked in the specified number of seconds, the system reloads. If you set the time value too low, it is possible for the system to reboot immediately after it loads.

Values: 15, 60, or 180

Default: 180 seconds (rounded up to 291 seconds by the Junos kernel)

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

Related Documentation

- *Disabling Junos OS Processes*

proxy-arp

Syntax proxy-arp;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.5.

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

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 Gratuitous ARP*

queue (Global Queues)

Syntax queue *queue-number class-name*;

Hierarchy Level [edit class-of-service forwarding-classes]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the output transmission queue to which to map all input from an associated forwarding class.

On M120, M320, MX Series, T Series routers and on EX Series switches, this statement enables you to configure up to eight forwarding classes with one-to-one mapping to output queues. If you want to configure up to 16 forwarding classes with multiple forwarding classes mapped to single output queues, include the **class** statement instead of the **queue** statement at the **[edit class-of-service forwarding-classes]** hierarchy level.

Options ***class-name***—Name of forwarding class.

queue-number—Output queue number.

Range: For M Series routers, 0 through 3. For M120, M320, MX Series, T Series routers and EX Series switches, 0 through 7. Some T Series router PICs are restricted to 0 through 3.

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

Related Documentation
- *Configuring Forwarding Classes*
- *class (Forwarding Classes)*

radius (Access Profile)

Syntax

```
radius {
  accounting-server [ ip-address ];
  attributes {
    exclude
      ...
    }
    ignore {
      framed-ip-netmask;
      input-filter;
      logical-system-routing-instance;
      output-filter;
    }
  }
  authentication-server [ ip-address ];
  options {
    accounting-session-id-format (decimal | description);
    calling-station-id-delimiter delimiter-character;
    calling-station-id-format {
      agent-circuit-id;
      agent-remote-id;
      interface-description;
      nas-identifier;
    }
    client-accounting-algorithm (direct | round-robin);
    client-authentication-algorithm (direct | round-robin);
    coa-dynamic-variable-validation;
    ethernet-port-type-virtual;
    interface-description-format {
      exclude-adapter;
      exclude-sub-interface;
    }
    ip-address-change-notify message;
    juniper-dsl-attributes;
    nas-identifier identifier-value;
    nas-port-extended-format {
      adapter-width width;
      ae-width width;
      port-width width;
      slot-width width;
      stacked-vlan-width width;
      vlan-width width;
      atm {
        adapter-width width;
        port-width width:
        slot-width width;
        vci-width width:
        vpi-width width;
      }
    }
    nas-port-id-delimiter delimiter-character;
    nas-port-id-format {
      agent-circuit-id;
```

```
            agent-remote-id;
            interface-description;
            nas-identifier;
        }
        nas-port-type {
            ethernet {
                port-type;
            }
        }
        revert-interval interval;
        vlan-nas-port-stacked-format;
    }
}
```

Hierarchy Level [edit access profile *profile-name*]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.

Description Configure the RADIUS parameters that the router uses for AAA authentication and accounting for subscribers.

The remaining statements are explained separately.

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

Related Documentation

- *Configuring RADIUS Server Parameters for Subscriber Access*
- *RADIUS Server Options for Subscriber Access*

radius-server

Syntax

```
radius-server server-address {
   accounting-port port-number;
   port port-number;
   retry attempts;
   routing-instance routing-instance-name;
   secret password;
   max-outstanding-requests value;
   source-address source-address;
   timeout seconds;
}
```

Hierarchy Level

[edit access],
[edit access profile *profile-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 RADIUS for subscriber access management, L2TP, or PPP.

To configure multiple RADIUS servers, include multiple **radius-server** statements. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.

Options

server-address—Address of the RADIUS authentication server.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Configuring RADIUS Authentication for L2TP*
- *Configuring the PPP Authentication Protocol*
- *Configuring RADIUS Authentication*
- *Configuring Authentication and Accounting Parameters for Subscriber Access*
- *Configuring an EX Series Switch to Use Junos Pulse Access Control Service for Network Access Control (CLI Procedure)*
- show network-access aaa statistics on page 178
- *clear network-access aaa statistics*

range (Address-Assignment Pools)

Syntax

```
range range-name {
   high upper-limit;
   low lower-limit;
   prefix-length prefix-length;
}
```

Hierarchy Level

[edit access address-assignment pool *pool-name* family (inet | inet6)]

Release Information

Statement introduced in Junos OS Release 9.0.
IPv6 support introduced in Junos OS Release 10.0.
Statement introduced in Junos OS Release 12.3 for EX Series switches.

Description

Configure a named range of IPv4 addresses or IPv6 prefixes, used within an address-assignment pool.

Options

high *upper-limit*—Upper limit of an address range or IPv6 prefix range.

low *lower-limit*—Lower limit of an address range or IPv6 prefix range.

prefix-length *prefix-length*—Assigned length of the IPv6 prefix.

range-name—Name assigned to the range of IPv4 addresses or IPv6 prefixes.

Required Privilege Level

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

Related Documentation

- *Address-Assignment Pools Overview*
- *Configuring Address-Assignment Pools*
- *Configuring a DHCP Server on EX Series Switches (CLI Procedure)*

router (Address-Assignment Pools)

Syntax router [*router-address*];

Hierarchy Level [edit access address-assignment pool *pool-name* family inet dhcp-attributes]

Release Information Statement introduced in Junos OS Release 9.0.

Description Specify one or more routers located on the client's subnet. This statement is the equivalent of DHCP option 3.

Options ***router-address***—IP address of one or more routers.

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

Related Documentation
- *Configuring Address-Assignment Pools*

scheduler-map (Dynamic Traffic Shaping)

Syntax scheduler-map (*map-name*);

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]

Release Information Statement introduced in Junos OS Release 9.3.
The **$junos-cos-scheduler-map** variable introduced in Junos OS Release 9.4.

Description Associate a scheduler map name with a traffic-control profile in a dynamic profile.

The scheduler map can be defined dynamically (at the **[edit dynamic-profiles *profile-name* class-of-service scheduler-maps]** hierarchy level) or statically (at the **[edit class-of-service scheduler-maps]** hierarchy level).

Options ***map-name***—Name of the scheduler map or the Junos predefined variable (**$junos-cos-scheduler-map**). When you specify the variable, the scheduler-map name is obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

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

Related Documentation
- *Guidelines for Configuring Dynamic CoS for Subscriber Access*
- *Configuring Traffic Scheduling and Shaping for Subscriber Access*
- output-traffic-control-profile on page 118

scheduler-maps (For Most Interface Types)

Syntax

```
scheduler-maps {
  map-name {
    forwarding-class class-name scheduler scheduler-name;
  }
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify a scheduler map name and associate it with the scheduler configuration and forwarding class.

Options *map-name*—Name of the scheduler map.

The remaining statements are explained separately.

See *Configuring Schedulers* and *Example: Configuring CoS for a PBB Network on MX Series Routers*.

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

schedulers (Class of Service)

Syntax

```
schedulers {
  scheduler-name {
    adjust-minimum rate;
    adjust-percent percentage;
    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 | medium-low | medium-high | high | none];
    excess-rate (percent percentage | proportion value);
    priority priority-level;
    shaping-rate (percent percentage | rate);
    transmit-rate (percent percentage | rate | remainder) <exact | rate-limit>;
  }
}
```

Hierarchy Level

[edit class-of-service]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Series routers.

Description

Specify the scheduler name and parameter values.

Options

scheduler-name—Name of the scheduler to be configured.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Schedulers Overview*
- *Default Schedulers Overview*
- *Configuring Schedulers*
- *Configuring a Scheduler*
- *Example: Configuring CoS for a PBB Network on MX Series Routers*

secret

Syntax secret *password*;

Hierarchy Level [edit access profile *profile-name* radius-server *server-address*],
[edit access radius-disconnect *client-address*],
[edit access radius-server *server-address*]

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 to use with the RADIUS server. The secret password used by the local router or switch must match that used by the server.

Options ***password***—Password to use; it can include spaces if the character string is enclosed in quotation marks.

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

Related Documentation

- *Configuring Authentication and Accounting Parameters for Subscriber Access*
- *Configuring Router or Switch Interaction with RADIUS Servers*
- *Example: Configuring CHAP Authentication with RADIUS*
- *Configuring RADIUS Authentication for L2TP*
- *Configuring the RADIUS Disconnect Server for L2TP*
- *Configuring an EX Series Switch to Use Junos Pulse Access Control Service for Network Access Control (CLI Procedure)*

services (System Services)

Syntax

```
services {
  dhcp { \* DHCP not supported on a DCF
    dhcp_services;
  }
  finger {
    connection-limit limit;
    rate-limit limit;
  }
  ftp {
    connection-limit limit;
    rate-limit limit;
  }
  service-deployment {
    servers address {
      port-number port-number;
    }
    source-address address;
  }
  ssh {
    connection-limit limit;
    protocol-version [v1 v2];
    rate-limit limit;
    root-login (allow | deny | deny-password);
  }
  telnet {
    connection-limit limit;
    rate-limit limit;
  }
  web-management {
    http {
      interfaces [ names ];
      port port;
    }
    https {
      interfaces [ names ];
      local-certificate name;
      port port;
    }
    session {
      idle-timeout [ minutes ];
      session-limit [ limit ];
    }
  }
  xnm-clear-text {
    connection-limit limit;
    rate-limit limit;
  }
  xnm-ssl {
    connection-limit limit;
    local-certificate name;
    rate-limit limit;
  }
```

```
}
```

Hierarchy Level [edit system]

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 router or switch so that users on remote systems can access the local router or switch through the DHCP server, finger, rlogin, SSH, telnet, Web management, Junos XML protocol clear-text, Junos XML protocol SSL, and network utilities or enable Junos OS to work with the Session and Resource Control (SRC) software.

The remaining statements are explained separately.

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

Related Documentation
- *Configuring clear-text or SSL Service for Junos XML Protocol Client Applications*
- *Configuring the Router or Interface to Act as a DHCP Server on J Series Services Routers*
- *Configuring the Junos OS to Work with SRC Software*

server (Dynamic PPPoE)

Syntax server;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" pppoe-options]

Release Information Statement introduced in Junos OS Release 10.1.

Description In a dynamic profile, configure the router to act as a PPPoE server, also known as a remote access concentrator, when a PPPoE logical interface is dynamically created.

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

Related Documentation
- *Configuring a Basic PPPoE Dynamic Profile*
- For information about creating static PPPoE interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*

shaping-rate (Applying to an Interface)

Syntax shaping-rate *rate*;

Hierarchy Level [edit class-of-service interfaces *interface-name*],
[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.
[edit class-of-service interfaces *interface-name*] hierarchy level added in Junos OS Release 7.5.
Statement introduced in Junos OS Release 13.2 on PTX Series Packet Transport Routers.

Description For logical interfaces on which you configure packet scheduling, configure traffic shaping by specifying the amount of bandwidth to be allocated to the logical interface.

For physical interfaces on IQ PICs and T4000 routers with Type 5 FPCs only, configure traffic shaping based on the rate-limited bandwidth of the total interface bandwidth.

Logical and physical interface traffic shaping rates are mutually exclusive. This means you can include the **shaping-rate** statement at the **[edit class-of-service interfaces *interface-name*]** hierarchy level or the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy level, but not both.

NOTE: For MX Series routers and for EX Series switches, the shaping rate value for the physical interface at the **[edit class-of-service interfaces *interface-name*]** hierarchy level must be a minimum of 160 Kbps. If the value is less than the sum of the logical interface guaranteed rates, the user is not allowed to apply the shaping rate to a physical interface.

For T4000 routers with Type 5 FPCs, the shaping rate value for the physical interface must be a minimum of 292 Kbps. The maximum value of **shaping-rate** is limited by the maximum transmission rate of the interface.

Alternatively, you can configure a shaping rate for a logical interface and oversubscribe the physical interface by including the **shaping-rate** statement at the **[edit class-of-service traffic-control-profiles]** hierarchy level. With this configuration approach, you can independently control the delay-buffer rate, as described in *Oversubscribing Interface Bandwidth*.

For FRF.15 and FRF.16 bundles on link services interfaces, only shaping rates based on percentage are supported.

Default If you do not include this statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy level, the default logical interface bandwidth is the average of unused bandwidth for the number of logical interfaces that require default bandwidth treatment. If you do not include this statement at the **[edit class-of-service interfaces *interface-name*]** hierarchy level, the default physical interface bandwidth is the

average of unused bandwidth for the number of physical interfaces that require default bandwidth treatment.

Options

rate—Peak rate, in bits per second (bps). 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).

Range: For logical interfaces, 1000 through 32,000,000,000 bps. For physical interfaces, 1000 through 160,000,000,000 bps.

NOTE: **For all MX Series and EX series interfaces, the rate can be from 65,535 through 160,000,000,000 bps.**

NOTE: **For T4000 physical interfaces, the rate can be from 1000 through 160,000,000,000 bps.**

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 Overview*
- *Configuring Virtual LAN Queuing and Shaping on PTX Series Packet Transport Routers*

shaping-rate (Dynamic Traffic Shaping and Scheduling)

Syntax shaping-rate (*rate* | *predefined-variable*) <burst-size *bytes* | $junos-cos-shaping-rate-burst>;

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*], [edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]

Release Information Statement introduced in Junos OS Release 9.2.
The **$junos-cos-shaping-rate** variable for traffic-control profiles introduced in Junos OS Release 9.4.
The **$junos-cos-scheduler-shaping-rate** variable for schedulers introduced in Junos OS Release 10.2.
Option **burst-size** introduced in Junos OS Release 11.4.

Description Configure a shaping rate for a logical interface or a scheduler. The sum of the shaping rates for all logical interfaces on the physical interface can exceed the physical interface bandwidth. This practice is known as oversubscription of the peak information rate (PIR).

Options *rate*—Peak rate in bits per second (bps). You can specify the value 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: 1000 through 160,000,000,000 bps

predefined-variable—One of the following Junos predefined variables. The variable is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

- **$junos-cos-shaping-rate**—Variable for the shaping rate that is specified for the logical interface. Use this variable at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]** hierarchy level.
- **$junos-cos-scheduler-shaping-rate**—Variable for the shaping rate that is specified for a scheduler. Use this variable at the **[edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]** hierarchy level.

burst-size *bytes*—(Optional) Maximum burst size, in bytes.

Range: 0 through 1,000,000,000

$junos-cos-shaping-rate-burst—(Optional) Variable for the burst-size that is specified for the shaping rate. Use this variable at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profile]** hierarchy level.

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

Related Documentation

- *Guidelines for Configuring Dynamic CoS for Subscriber Access*
- *Configuring Traffic Scheduling and Shaping for Subscriber Access*
- output-traffic-control-profile on page 118

source-address

Syntax source-address *source-address*;

Hierarchy Level [edit access radius-server *server-address*],
[edit access profile *profile-name* radius-server *server-address*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure a source address for each configured RADIUS server. Each RADIUS request sent to a RADIUS server uses the specified source address.

Options ***source-address***—Valid IPv4 address configured on one of the router or switch interfaces. On M Series routers only, the source address can be an IPv6 address and the UDP source port is 514.

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

Related Documentation

- *Configuring Router or Switch Interaction with RADIUS Servers*
- *Configuring Authentication and Accounting Parameters for Subscriber Access*
- *Example: Configuring CHAP Authentication with RADIUS*
- *Configuring RADIUS Authentication for L2TP*

statistics (Access Profile)

Syntax statistics (time | volume-time);

Hierarchy Level [edit access profile *profile-name* accounting]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.
Option **volume-time** introduced in Junos OS Release 9.4.

Description Configure the router or switch to collect time statistics, or both volume and time statistics, for the sessions being managed by AAA.

Options **time**—Collect uptime statistics only.

volume-time—Collect both volume and uptime statistics. This option is not available for Mobile IP.

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

Related Documentation

- *Mobile IP Home Agent Elements and Behavior*
- *Configuring Authentication and Accounting Parameters for Subscriber Access*

stacked-vlan-ranges

Syntax

```
stacked-vlan-ranges {
  access-profile profile-name;
  authentication {
    password password-string;
    username-include {
      circuit-type;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-name;
      mac-address;
      option-18
      option-37
      option-82;
      radius-realm radius-realm-string;
      user-prefix user-prefix-string;
    }
  }
  dynamic-profile profile-name {
    accept (any | dhcp-v4 | inet);
    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.

Required Privilege Level

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

Related Documentation

- *Configuring Stacked VLAN Ranges for Use with Stacked VLAN Dynamic Profiles*
- *Configuring Dynamic Mixed VLAN Ranges*

system

Syntax system { ... }

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure system management properties.

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

Related Documentation

- *System Management Configuration Statements*

term (Firewall Filter)

Syntax

```
term term-name {
  from {
    match-conditions;
    ip-version ipv4 {
      match-conditions-mpls-ipv4-address;
        protocol (tcp | udp) {
          match conditions-mpls-ipv4-port;
        }
    }
  }
  then {
    actions;
  }
}
```

Hierarchy Level

[edit firewall family *family-name* filter *filter-name*],
[edit firewall family *family-name* service-filter *filter-name*],
[edit firewall family *family-name* simple-filter *filter-name*],
[edit logical-systems *logical-system-name* firewall family *family-name* filter *filter-name*],
[edit logical-systems *logical-system-name* firewall family *family-name* service-filter *filter-name*],
[edit logical-systems *logical-system-name* firewall family *family-name* simple-filter *filter-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
filter option introduced in Junos OS Release 7.6.
Logical systems support introduced in Junos OS Release 9.3.
ip-version ipv4 support introduced in Junos OS Release 10.1.

Description

Define a firewall filter term.

Options

actions—(Optional) Actions to perform on the packet if conditions match. You can specify one *terminating action* supported for the specified filter type. If you do not specify a terminating action, the packets that match the conditions in the **from** statement are accepted by default. As an option, you can specify one or more *nonterminating actions* supported for the specified filter type.

filter-name—(Optional) For **family *family-name* filter *filter-name*** only, reference another standard stateless firewall filter from within this term.

from—(Optional) Match packet fields to values. If not included, all packets are considered to match and the actions and action modifiers in the **then** statement are taken.

match-conditions—One or more conditions to use to make a match on a packet.

match-conditions-mpls-ipv4-address—(MPLS-tagged IPv4 traffic only) One or more IP address match conditions to match on the IPv4 packet header. Supports network-based service in a core network with IPv4 packets as an inner payload of an MPLS packet with labels stacked up to five deep.

match-conditions-mpls-ipv4-port—(MPLS-tagged IPv4 traffic only) One or more UDP or TCP port match conditions to use to match a packet in an MPLS flow. Supports network-based service in a core network with IPv4 packets as an inner payload of an MPLS packet with labels stacked up to five deep.

term-name—Name that identifies the term. The name can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks (“ ”).

then—(Optional) Actions to take on matching packets. If not included and a packet matches all the conditions in the **from** statement, the packet is accepted.

Required Privilege Level

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

Related Documentation

- *Guidelines for Configuring Firewall Filters*
- *Configuring Multifield Classifiers*
- *Guidelines for Configuring and Applying Firewall Filters in Logical Systems*

traffic-control-profiles (Dynamic CoS Definition)

Syntax

```
traffic-control-profiles profile-name {
   adjust-minimum rate;
   delay-buffer-rate (percent percentage | rate);
   excess-rate (percent percentage | proportion value | percent $junos-cos-excess-rate);
   excess-rate-high (percent percentage | proportion value);
   excess-rate-low (percent percentage | proportion value);
   guaranteed-rate (percent percentage | rate) <burst-size  bytes>;
   overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
   scheduler-map map-name;
   shaping-rate (percent percentage | rate | predefined-variable) <burst-size  bytes>;
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* class-of-service]

Release Information

Statement introduced in Junos OS Release 9.2.

Description

Configure traffic shaping and scheduling profiles.

Options

profile-name—Name of the traffic-control profile.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Guidelines for Configuring Dynamic CoS for Subscriber Access*
- *Configuring Traffic Scheduling and Shaping for Subscriber Access*
- *Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers*
- output-traffic-control-profile on page 118

transmit-rate (Schedulers)

Syntax transmit-rate (*rate* | percent *percentage* | remainder) <exact | rate-limit>;

Hierarchy Level [edit class-of-service schedulers *scheduler-name*]

Release Information Statement introduced before Junos OS Release 7.4.
rate-limit option introduced in Junos OS Release 8.3. Applied to the Multiservices PICs in Junos OS Release 9.4.
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 Routers.

Description Specify the transmit rate or percentage for a scheduler.

Default If you do not include this statement, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent, respectively.

Options **exact**—(Optional) Enforce the exact transmission rate. Under sustained congestion, a rate-controlled queue that goes into negative credit fills up and eventually drops packets. This value should never exceed the rate-controlled amount. For PTX Series Packet Transport Routers, this option is allowed only on the non-strict-high (high, medium-high, medium-low, or low) queues.

percent *percentage*—Percentage of transmission capacity. A percentage of zero drops all packets in the queue.

Range: 0 through 100 percent for M, MX and T Series routers and EX Series switches; 1 through 100 percent for PTX Series Packet Transport Routers; 0 through 200 percent for the SONET/SDH OC48/STM16 IQE PIC

NOTE:

- On M Series Multiservice Edge Routers, for interfaces configured on 4-port E1 and 4-port T1 PICs only, you can configure a *percentage* value only from 11 through 100. These two PICs do not support transmission rates less than 11 percent.
- The configuration of the **transmit-rate percent 0 exact** statement at the [**edit class-of-service schedules** *scheduler-name*] hierarchy is ineffective on T4000 routers with Type 5 FPC.
- On MIC and MPC interfaces on MX Series routers, when the transmit rate is configured as a percentage and **exact** or **rate-limit** is enabled on a queue, the shaping rate of the parent node is used to compute the transmit rate. If **exact** or **rate-limit** is not configured, the guaranteed rate of the parent node is used to compute the transmit rate.

rate—Transmission rate, in bps. 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).

Range: 3200 through 160,000,000,000 bps

NOTE: For all MX Series interfaces, the rate can be from 65,535 through 160,000,000,000 bps.

rate-limit—(Optional) Limit the transmission rate to the rate-controlled amount. In contrast to the **exact** option, the scheduler with the **rate-limit** option shares unused bandwidth above the rate-controlled amount.

NOTE: For PTX Series Packet Transport Routers, this option is allowed only on the strict-high queue. We recommend that you configure rate limit on strict-high queues because the other queues may not meet their guaranteed bandwidths.

NOTE: The configuration of the **rate-limit** statement is supported on T4000 routers only with a Type 5 FPC.

remainder—Use the remaining rate available.

Required Privilege Level

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

Related Documentation

- *Configuring Schedulers*
- *Configuring Scheduler Transmission Rate*
- *Example: Configuring CoS for a PBB Network on MX Series Routers*

underlying-interface (Dynamic PPPoE)

Syntax underlying-interface *interface-name*;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces pp0 unit "$junos-interface-unit" pppoe-options]

Release Information Statement introduced in Junos OS Release 10.1.

Description In a dynamic profile, configure the underlying interface on which the router creates the dynamic PPPoE logical interface.

Options ***interface-name***—Variable used to specify the name of the underlying interface on which the PPPoE logical interface is dynamically created. In the **underlying-interface *interface-name*** statement for dynamic PPPoE logical interfaces, you must use the predefined variable **$junos-underlying-interface** in place of ***interface-name***. When the router creates the dynamic PPPoE interface, the **$junos-underlying-interface** predefined variable is dynamically replaced with the name of the underlying interface supplied by the network when the subscriber logs in.

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

Related Documentation

- *Configuring a Basic PPPoE Dynamic Profile*
- For information about creating static PPPoE interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*

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);
  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;
  }
  description text;
  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;
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;
   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 {
      access-profile name;
      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;
}
pppoe-underlying-options {
   access-concentrator name;
   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 ...
   }
   bridge-domain-type (bvlan | svlan);
   bundle interface-name;
   core-facing;
   demux-destination {
      destination-prefix;
   }
   demux-source {
      source-prefix;
   }
   duplicate-protection;
   dynamic-profile profile-name;
   filter {
      group filter-group-number;
      input filter-name;
      input-list [ filter-names ];
      output filter-name;
      output-list [ filter-names ];
   }
   interface-mode (access | trunk);
   ipsec-sa sa-name;
   isid-list all-service-groups;
   keep-address-and-control;
   mac-validate (loose | strict);
   max-sessions number;
```

```
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];
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;
              }
            }
          }
        }
      }
```

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 only. 0 through 16,385 for all other static interface types.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Configuring Logical Interface Properties*
- *Example: Configuring E-LINE and E-LAN Services for a PBB Network on MX Series Routers*
- *Junos OS Services Interfaces Library for Routing Devices*

unit

Syntax

```
unit logical-unit-number {
  classifiers {
    type (classifier-name | default) family (mpls | all);
  }
  forwarding-class class-name;
  fragmentation-map map-name;
  input-traffic-control-profile profile-name shared-instance instance-name;
  output-traffic-control-profile profile-name shared-instance instance-name;
  per-session-scheduler;
  rewrite-rules {
    dscp (rewrite-name | default);
    dscp-ipv6 (rewrite-name | default);
    exp (rewrite-name | default)protocol protocol-types;
    exp-push-push-push default;
    exp-swap-push-push default;
    ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
    inet-precedence (rewrite-name | default);
  }
  scheduler-map map-name;
  shaping-rate rate;
}
```

Hierarchy Level

[edit class-of-service interfaces *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 16,384

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Overview of BA Classifier Types*
- *Configuring Rewrite Rules*

unit (Dynamic Profiles Standard Interface)

Syntax

```
unit logical-unit-number {
  auto-configure {
    agent-circuit-identifier {
      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 {
    access-concentrator name;
    address address;
    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;
      }
      output filter-name {
        precedence precedence;
      }
    }
    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;
      }
      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>;
        filter {
          input filter-name;
          output filter-name;
        }
        keepalives {
          interval seconds;
          }
        ppp-options {
          chap;
          pap;
        }
        vlan-id number;
        vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
      }
    }
```

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 Junos OS 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 agent circuit identifier information.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- *Configuring 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 Components Overview*

unit (Dynamic Traffic Shaping)

Syntax

```
unit logical-unit-number {
  classifiers {
    type (classifier-name | default);
  }
  output-traffic-control-profile (profile-name | $junos-cos-traffic-control-profile);
  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);
    }
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* class-of-service interfaces *interface-name*],
[edit dynamic-profiles *profile-name* interfaces interface-set *interface-set-name* interface *interface-name*]

Release Information

Statement introduced in Junos OS Release 9.2.
Support at the **[edit dynamic-profiles *profile-name* class-of-service interfaces interface-set *interface-set-name*]** hierarchy level introduced in Junos OS Release 10.4.

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—One of the following options:

- **$junos-underlying-interface-unit**—For static VLANs, the unit number variable. The static unit number variable is dynamically replaced with the client unit number when the client session begins. The client unit number is specified by the DHCP when it accesses the subscriber network.
- **$junos-interface-unit**—For dynamic demux and dynamic PPPoE interfaces, 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 or PPP when it accesses the subscriber network.
- *value*—Specific unit number of the interface you want to assign to the dynamic-profile

Range: 0 through 16385. For demux and PPPoE interfaces, the unit numbers can range from 0 through 1,073,741,823.

The remaining statements are explained separately. The **classifiers**, **output-traffic-control-profile**, and **rewrite-rules** statements are not supported for interface sets.

Required Privilege Level

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

Related Documentation

- *Guidelines for Configuring Dynamic CoS for Subscriber Access*
- *Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile*
- *Configuring an Interface Set of Subscribers in a Dynamic Profile*

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*
- address on page 59
- *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.
$junos-preferred-source-address variable support added in Junos OS Release 9.6.
Support for the **$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. Use the **$junos-loopback-interface** dynamic variable to dynamically apply a loopback interface. The loopback interface used is based on the routing instance of the subscriber. The specified interface must have a logical unit number and a configured IP address, and must not be an unnumbered interface.

preferred-source-address *address*—(Optional) Secondary IP address of the donor loopback interface. Use the **$junos-preferred-source-address** dynamic variable to dynamically apply a preferred source address to the unnumbered Ethernet interface. When you use the dynamic variable, the address that is selected resides in the same network as the IP address of the subscriber. 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

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* in *Junos OS Network Interfaces Library for Routing Devices*.
- *Junos OS Network Interfaces Library for Routing Devices*

update-interval

Syntax update-interval *minutes*;

Hierarchy Level [edit access profile *profile-name* accounting]

Release Information Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 9.1 for EX Series switches.

Description Enable interim accounting updates and configure the amount of time that the router or switch waits before sending a new accounting update.

Default No updates

Options ***minutes***—Amount of time between updates, in minutes. All values are rounded up to the next higher multiple of 10. For example, the values 811 through 819 are all accepted by the CLI, but are all rounded up to 820.

Range: 10 through 1440 minutes

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

Related Documentation

- *Configuring Authentication and Accounting Parameters for Subscriber Access*

username-include (DHCP Local Server)

Syntax

```
username-include {
   circuit-type;
   client-id;
   delimiter delimiter-character;
   domain-name domain-name-string;
   interface-name ;
   logical-system-name;
   mac-address;
   option-60;
   option-82 <circuit-id> <remote-id>;
   relay-agent-interface-id;
   relay-agent-remote-id;
   relay-agent-subscriber-id;
   routing-instance-name;
   user-prefix user-prefix-string;
}
```

Hierarchy Level

[edit system services dhcp-local-server authentication],
[edit system services dhcp-local-server dhcpv6 authentication],
[edit system services dhcp-local-server dhcpv6 group *group-name* authentication],
[edit system services dhcp-local-server group *group-name* authentication],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server ...],
[edit logical-systems *logical-system-name* system services dhcp-local-server ...],
[edit routing-instances *routing-instance-name* system services dhcp-local-server ...]

Release Information

Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure the username that the router or switch 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 (or switch) accesses the local authentication service only and does not use external authentication services, such as RADIUS.

The statements are explained separately. The *option-60* and *option-82* statements are not supported in the DHCPv6 hierarchy levels. The *client-id*, *relay-agent-interface-id*, *relay-agent-remote-id* and *relay-agent-subscriber-id* statements are supported in the DHCPv6 hierarchy levels only.

Required Privilege Level

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

Related Documentation

- *Using External AAA Authentication Services with DHCP*
- *Creating Unique Usernames for DHCP Clients*

user-prefix (DHCP Local Server)

Syntax

user-prefix *user-prefix-string*;

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server authentication username-include],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 authentication username-include],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 group *group-name* authentication username-include],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* authentication username-include],
[edit logical-systems *logical-system-name* system services dhcp-local-server authentication username-include],
[edit logical-systems *logical-system-name* system services dhcp-local-server dhcpv6 authentication username-include],
[edit logical-systems *logical-system-name* system services dhcp-local-server dhcpv6 group *group-name* authentication username-include],
[edit logical-systems *logical-system-name* system services dhcp-local-server group *group-name* authentication username-include],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server authentication username-include],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 authentication username-include],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 group *group-name* authentication username-include],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* authentication username-include],
[edit routing-instances *routing-instance-name* system services dhcp-local-server authentication username-include],
[edit routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 authentication username-include],
[edit routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 group *group-name* authentication username-include],
[edit routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* authentication username-include],
[edit system services dhcp-local-server authentication username-include],
[edit system services dhcp-local-server dhcpv6 authentication username-include],
[edit system services dhcp-local-server dhcpv6 group *group-name* authentication username-include],
[edit system services dhcp-local-server group *group-name* authentication username-include]

Release Information

Statement introduced in Junos OS Release 9.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Specify the user prefix that is concatenated with the username during the subscriber authentication or DHCP client authentication process.

Options

user-prefix-string—User prefix string.

Required Privilege Level

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

Related Documentation

- *Using External AAA Authentication Services with DHCP*

vlan-ranges

Syntax

```
vlan-ranges {
  access-profile profile-name;
  authentication {
    password password-string;
    username-include {
      circuit-type;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-name;
      mac-address;
      option-82 <circuit-id> <remote-id>;
      radius-realm radius-realm-string;
      user-prefix user-prefix-string;
    }
  }
  dynamic-profile profile-name {
    accept (any | dhcp-v4 | inet);
    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.

Required Privilege Level

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

Related Documentation

- *Configuring Single-Level VLAN Ranges for Use with VLAN Dynamic Profiles*
- *Configuring Dynamic Mixed VLAN Ranges*

vlan-tags (Stacked VLAN Tags)

Syntax vlan-tags inner *tpid.vlan-id* 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 For Gigabit Ethernet IQ and IQE interfaces only, bind TPIDs and 802.1Q VLAN tag IDs to a logical interface.

NOTE: The inner-range *vid1–vid2* option is supported on MX Series with IQE PICs only.

Options inner *tpid.vlan-id*—A TPID and a valid VLAN identifier.

Range: (most routers) For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

Range: (PTX Series) For VLAN ID, 0 through 4094.

inner-range *vid1–vid2*—For MX Series routers with Enhanced IQ (IQE) PICs only; specify a range of VLAN IDs where vid1 is the start of the range and vid2 is the end of the range.

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.

Range: (PTX Series) For VLAN ID, 0 through 511 for normal interfaces, and 512 through 4094 for VLAN CCC interfaces.

NOTE: Configuring inner-range with the entire vlan-id range consumes system resources and is not a 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 (1–4094), it 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*

vlan-tags

Syntax vlan-tags outer [*tpid*].*vlan-id* [inner [*tpid*].*vlan-id*];

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.5.
VLAN demux interface support introduced in Junos OS Release 10.2.

Description For Gigabit Ethernet IQ and IQE interfaces only, binds TPIDs and 802.1Q VLAN tag IDs to a logical interface. You must include the **stacked-vlan-tagging** statement at the **[edit interfaces *interface-name*]** hierarchy level.

NOTE: The **inner-range *vid1–vid2*** option is supported on MX Series routers with IQE PICs only.

Options

inner [*tpid*].*vlan-id*—A TPID (optional) and a valid VLAN identifier in the format *tpid.vlan-id*. When used in the **dynamic-profiles** hierarchy, specify the **$junos-vlan-id** predefined variable to dynamically obtain the VLAN ID.

Range: For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

outer [*tpid*].*vlan-id*—A TPID (optional) and a valid VLAN identifier in the format *tpid.vlan-id*. When used in the **dynamic-profiles** hierarchy, specify the **$junos-stacked-vlan-id** predefined variable.

Range: 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.

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*
- *stacked-vlan-tagging*

PART 3

Administration

CHAPTER 7

Subscriber Management AAA and DHCP CLI Commands

show network-access aaa statistics

Syntax

show network-access aaa statistics
<accounting>
<address-assignment (client | pool *pool-name*)>
<dynamic-requests>
<radius>

Release Information

Command introduced in Junos OS Release 9.1.
Option **address-assignment** introduced in Junos OS Release 10.0.
Option **radius** introduced in Junos OS Release 11.4.

Description

Display AAA accounting, address-assignment, dynamic request statistics, and RADIUS settings and statistics.

Options

accounting—(Optional) Display AAA accounting statistics.

address-assignment (client | pool *pool-name*)—(Optional) Display AAA address-assignment client and pool statistics.

dynamic-requests—(Optional) Display AAA dynamic requests.

radius— (Optional) Display RADIUS settings and statistics.

Required Privilege Level

view

Related Documentation

- *Verifying and Managing Subscriber AAA Information*

List of Sample Output

show network-access aaa statistics accounting on page 180
show network-access aaa statistics address-assignment client on page 180
show network-access aaa statistics address-assignment pool on page 180
show network-access aaa statistics dynamic-requests on page 180
show network-access aaa statistics radius on page 180

Output Fields

Table 6 on page 178 lists the output fields for the **show network-access aaa statistics** command. Output fields are listed in the approximate order in which they appear.

Table 6: show network-access aaa statistics Output Fields

Field Name	Field Description
Requests received	• Number of accounting requests generated by the AAA framework. • Number of dynamic requests received from the external server.
Accounting Response failures	Number of accounting requests not acknowledged (NAK) by the accounting server.
Accounting Response Success	Number of accounting requests acknowledged by the accounting server.

Table 6: show network-access aaa statistics Output Fields *(continued)*

Field Name	Field Description
Requests timedout	Number of accounting requests to the accounting server that timed out.
Client	Client type; for example, DHCP, Mobile IP, PPP.
Out of Memory	Number of times an address was not given to the client due to memory issues.
No Matches	Number of times there were no network matches for the pool.
Pool Name	Name of the address-assignment pool for this client.
Out of Addresses	Number of times there were no available addresses in the pool.
Address total	Number of addresses in the pool.
Addresses in use	Number of addresses in use.
Address Usage (percent)	Percentage of total addresses in use.
processed successfully	Number of dynamic requests processed successfully by the AAA framework.
errors during processing	Number of dynamic requests that resulted in processing errors by the AAA framework.
Link Name	Name of the secondary address-assignment pool to which the primary pool is linked.
Pool Usage	Percentage of allocated addresses in the specified address pool.
silently dropped	Number of dynamic requests dropped by the AAA framework due to multiple back-to-back or duplicate requests.
RADIUS Server	IP address of the RADIUS server to which the router is sending requests.
Profile	Name of the RADIUS profile associated with the RADIUS server. A RADIUS server can be associated with more than one RADIUS profile.
Configured	Configured maximum number of outstanding requests from the router to the RADIUS server for a specific profile. An outstanding request is a request to which the RADIUS server has not yet responded. The range of values is 0 through 2000 outstanding requests. The default value is 1000.
Current	Current number of outstanding requests from the router to the RADIUS server for a specific profile. An outstanding request is a request to which the RADIUS server has not yet responded.
Peak	Highest number of outstanding requests from the router to the RADIUS server for a specific profile at any point in time since the router was started or since the counter was last cleared. NOTE: If the value of this field is equal to the value of the **Configured** field, you may want to increase the value of the **Configured** field.

Table 6: show network-access aaa statistics Output Fields *(continued)*

Field Name	Field Description
Exceeded	Number of times that the router attempted to send requests to the RADIUS server in excess of the configured maximum value for a specific profile. NOTE: If the value of this field is nonzero, you may want to increase the value of the **Configured** field.

Sample Output

show network-access aaa statistics accounting

```
user@host> show network-access aaa statistics accounting
Accounting module statistics
  Requests received: 0
  Accounting Response failures: 0
  Accounting Response Success: 0
  Requests timedout: 0
```

show network-access aaa statistics address-assignment client

```
user@host> show network-access aaa statistics address-assignment client
Address-assignment statistics
 Client: jdhcpd
 Out of Memory: 0
 No Matches: 2
```

show network-access aaa statistics address-assignment pool

```
user@host> show network-access aaa statistics address-assignment pool isp_1
Address-assignment statistics
 Pool Name: isp_1
 Pool Name: (all pools in chain)
 Out of Memory: 0
 Out of Addresses: 9
 Address total: 47
 Addresses in use: 47
 Address Usage (percent): 100
```

show network-access aaa statistics dynamic-requests

```
user@host> show network-access aaa statistics dynamic-requests
requests received: 0
processed successfully: 0
errors during processing: 0
silently dropped: 0
```

show network-access aaa statistics radius

```
user@host> show network-access aaa statistics radius
Outstanding Requests
RADIUS Server     Profile      Configured   Current   Peak   Exceeded
172.28.32.239     prof1        1000         0         1000   14
                  prof2        500          17        432    0
171.27.82.211     myprof       200          0         200    27
12.1.11.254     pppoe-auth     111          0         1      0
```

show network-access aaa statistics authentication

Syntax

show network-access aaa statistics authentication
<detail>

Release Information

Command introduced in Junos OS Release 9.1.
Option **detail** introduced in Junos OS Release 12.1.

Description

Display AAA authentication statistics.

Options

detail—(Optional) Displays detailed information about authentication.

Required Privilege Level

view

Related Documentation

- *Verifying and Managing Subscriber AAA Information*

List of Sample Output

show network-access aaa statistics authentication on page 183
show network-access aaa statistics authentication detail on page 183

Output Fields

Table 7 on page 181 lists the output fields for the **show network-access aaa statistics authentication** command. Output fields are listed in the approximate order in which they appear.

Table 7: show network-access aaa statistics authentication Output Fields

Field Name	Field Description	Level of Output
Requests received	Number of authentication requests received from clients.	All levels
Multistack requests	Number of authentication requests for dual-stack subscribers.	All levels
Accepts	Number of authentication requests accepted by the authentication server.	All levels
Rejects	Number of authentication requests rejected by the authentication server.	All levels
Challenges	Number of authentication requests challenged by the authentication server.	All levels
Requests timed out	Number of authentication requests that timed out.	All levels
RADIUS authentication failures	Number of RADIUS authentication requests that have failed.	Detail
Queue request deleted	Number of queue requests that have been deleted.	Detail

Table 7: show network-access aaa statistics authentication Output Fields *(continued)*

Field Name	Field Description	Level of Output
Malformed reply	Number of malformed replies received from the RADIUS authentication server.	Detail
No server configured	Number of authentication requests that failed because no authentication server is configured.	Detail
Access Profile configuration not found	Number of authentication requests that failed because no access profile is configured.	Detail
Unable to create client record	Number of times that the router is unable to create the client record for the authentication request.	Detail
Unable to create client request	Number of times that the router is unable to create the client request for the authentication request.	Detail
Unable to build authentication request	Number of times that the router is unable to build the authentication request.	Detail
No server found	Number of requests to the authentication server that have timed out; the server is then considered to be down.	Detail
Unable to create handle	Number of authentication requests that have failed because of an internal allocation failure.	Detail
Unable to queue request	Number of times the router was unable to queue the request to the authentication server.	Detail
Invalid credentials	Number of times the router did not have proper authorization to access the authentication server.	Detail
Malformed request	Number of times the router request to the authentication server is malformed.	Detail
License unavailable	Number of times the router did not have a license to access the authentication server.	Detail
Redirect requested	Number of authentication requests that have been redirected based on routing instance.	Detail
Internal failure	Number of internal failures.	Detail
Local authentication failures	Number of times local authentication failed.	Detail
LDAP lookup failures	Number of times the LDAP lookup operation failed.	Detail

Sample Output

show network-access aaa statistics authentication

```
user@host> show network-access aaa statistics authentication
Authentication module statistics
  Requests received: 2118
  Multistack requests: 0
  Accepts: 261
  Rejects: 975
  Challenges: 0
  Requests timed out: 882
```

show network-access aaa statistics authentication detail

```
user@host> show network-access aaa statistics authentication detail
Authentication module statistics
  Requests received: 2118
  Multistack requests: 0
  Accepts: 261
  Rejects: 975
    RADIUS authentication failures: 975
      Queue request deleted: 0
      Malformed reply: 0
      No server configured: 0
      Access Profile configuration not found: 0
      Unable to create client record: 0
      Unable to create client request: 0
      Unable to build authentication request: 0
      No server found: 975
      Unable to create handle: 0
      Unable to queue request: 0
      Invalid credentials: 0
      Malformed request: 0
      License unavailable: 0
      Redirect requested: 0
      Internal failure: 0
    Local authentication failures: 0
    LDAP lookup failures: 0
  Challenges: 0
  Requests timed out: 882
```

show network-access aaa subscribers

Syntax

show network-access aaa subscribers
<logical-system *logical-system-name*>
<routing-instance *routing-instance-name*>
<statistics>
<username>

Release Information

Command introduced in Junos OS Release 9.1.

Description

Display subscriber-specific AAA statistics.

Options

logical-system *logical-system-name*—(Optional) List subscribers in the specific logical system.

routing-instance *routing-instance-name*—(Optional) List subscribers for the specific routing instance. If you do not specify a routing instance name, the default routing instance is assumed.

statistics—(Optional) Display statistics for the subscriber events.

username—(Optional) Display information for the specified subscriber.

Required Privilege Level

view

Related Documentation

- *Verifying and Managing Subscriber AAA Information*

List of Sample Output

show network-access aaa subscribers logical-system on page 185
show network-access aaa subscribers logical-system routing-instance on page 185
show network-access aaa subscribers statistics username on page 186
show network-access aaa subscribers username on page 186

Output Fields

Table 8 on page 184 lists the output fields for the **show network-access aaa subscribers** command. Output fields are listed in the approximate order in which they appear.

Table 8: show network-access aaa subscribers Output Fields

Field Name	Field Description
Challenge requests	Number of authentication requests challenged by the authentication server for this subscriber.
Challenge responses	Number of challenge responses sent by the subscriber to the authentication server.
START sent successfully	Number of accounting start requests generated by the AAA framework for this subscriber.
START send failures	Number of accounting start requests that failed to make it to the accounting server for this subscriber.
START ack received	Number of accounting start requests acknowledged by the accounting server for this subscriber.

Table 8: show network-access aaa subscribers Output Fields *(continued)*

Field Name	Field Description
INTERIM sent successfully	Number of accounting interim requests generated by the AAA framework for this subscriber.
INTERIM send failures	Number of accounting interim requests that failed to make it to the accounting server for this subscriber.
INTERIM ack received	Number of accounting interim requests acknowledged by the accounting server for this subscriber.
Requests received	Number of reauthentication requests received by the authentication server.
Successful responses	Number of successful reauthentication requests granted by the authentication server.
Aborts handled	Number of reauthentication requests aborted by the authentication server.
Service name	Name of the subscriber service.
Creation requests	Number of requests to create the service.
Deletion requests	Number of requests to delete the service.
Request timeouts	Number of times the service request was timed out.
Client type	Type of client; for example, DHCP, Mobile IP, PPP.
Session-ID	ID of the subscriber session.
Session uptime	How long the session has been up, in *HH:MM:SS*.
Accounting	Status of accounting, and type of accounting if accounting is on.

Sample Output

show network-access aaa subscribers logical-system

```
user@host> show network-access aaa subscribers logical-system
Username            Client type    Logical system/Routing instance
cbenson@addr.net    ppp            default
00010e020304.1231   dhcp           isp-bos-metro-12:isp-cmbrg-12
conley@isp3.com     dhcp           default:isp-gtown-r3-00
0020df980102.2334   dhcp           isp-bos-metro-16:isp-cmbrg-12
```

show network-access aaa subscribers logical-system routing-instance

```
user@host> show network-access aaa subscribers logical-system isp-bos-metro-16
routing-instance isp-cmbrg-12-32
Username            Client type   Logical system/Routing instance
00010e020304.1231   dhcp          isp-bos-metro-12:isp-cmbrg-12
conley@isp3.com     dhcp          default:isp-gtown-r3-00
0020df980102.2334   dhcp          isp-bos-metro-16:isp-cmbrg-12
```

show network-access aaa subscribers statistics username

```
user@host> show network-access aaa subscribers statistics username 00010e020304.1231
Authentication statistics
   Challenge requests: 0
   Challenge responses: 0
 Accounting statistics
   START sent successfully: 1
   START send failures: 0
   START ack received: 1
   INTERIM sent successfully: 0
   INTERIM send failures: 0
   INTERIM ack received: 0
 Re-authentication statistics
   Requests received: 0
   Sucessfull responses: 0
   Aborts handled: 0
 Service statistics
   Service name: filter-serv
   Creation requests: 1
   Deletion requests: 0
   Request timeouts: 0
   Service name: filter-serv2
   Creation requests: 144
   Deletion requests: 0
   Request timeouts: 144
```

show network-access aaa subscribers username

```
user@host> show network-access aaa subscribers username fred@isp5.net
Logical system/Routing instance   Client type  Session-ID  Session uptime
Accounting
isp-bos-metro-16:isp-cmbrg-12     dhcp         7           01:12:56
on/volume
Service name       Service type      Quota          Accounting
I-Cast             volume            1200 Mbps      on/volume+time
Voip                                                on/volume
GamingBurst        time              6000 secs      on/volume
```

show network-access address-assignment pool

Syntax show network-access address-assignment pool *pool-name*
<logical-system *logical-system-name*>
<routing-instance *routing-instance-name*>

Release Information Command introduced in Junos OS Release 9.0.

Description Display state information for each address-assignment pool.

Options **none**—Display information about clients that have obtained addresses from the address-assignment pool.

pool *pool-name*—Display information about the specified address-assignment pool.

logical-system *logical-system-name*—(Optional) Perform this operation on the specified logical system.

routing-instance *routing-instance-name*—(Optional) Perform this operation on the specified routing instance.

Required Privilege Level view and system

List of Sample Output show network-access address-assignment pool on page 187

Output Fields Table 9 on page 187 lists the output fields for the **show address-assignment pool** command. Output fields are listed in the approximate order in which they appear.

Table 9: show network-access address-assignment pool Output Fields

Field Name	Field Description
IP address	IP address of the client.
Hardware address	MAC address of the client.
Type	Type of client.

Sample Output

show network-access address-assignment pool

```
user@host> show network-access address-assignment pool sunnywest logical-system ls1
routing-instance routinst2
IP address        Hardware address        Type
192.168.2.1       00:05:1b:00:b9:01       DHCP
192.168.2.2       00:05:1b:00:b9:02       DHCP
192.168.2.3       00:05:1b:00:b9:03       DHCP
192.168.2.4       00:05:1b:00:b9:04       DHCP
```

CHAPTER 8

Subscriber Management DHCP Local Server CLI Commands

show dhcp server binding

Syntax

show dhcp server binding
<*address*>
<**brief | detail | summary**>
<**interface** *interface-name*>
<*interfaces-vlan*>
<*interfaces-wildcard*>
<**logical-system** *logical-system-name*>
<**routing-instance** *routing-instance-name*>

Release Information

Command introduced in Junos OS Release 9.0.
Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

Description

Display the address bindings in the client table on the extended Dynamic Host Configuration Protocol (DHCP) local server.

Options

address—(Optional) Display DHCP binding information for a specific client identified by one of the following entries:

- *ip-address*—The specified IP address.
- *mac-address*—The specified MAC address.
- *session-id*—The specified session ID.

brief | detail | summary—(Optional) Display the specified level of output about active client bindings. The default is **brief**, which produces the same output as **show dhcp server binding**.

interface ***interface-name***—(Optional) Display information about active client bindings on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.

interfaces-vlan—(Optional) Show the binding state information on the interface VLAN ID and S-VLAN ID.

interfaces-wildcard—(Optional) The set of interfaces on which to show the binding state information. This option supports the use of the wildcard character (*).

logical-system ***logical-system-name***—(Optional) Display information about active client bindings for DHCP clients on the specified logical system.

routing-instance ***routing-instance-name***—(Optional) Display information about active client bindings for DHCP clients on the specified routing instance.

Required Privilege Level

view

Related Documentation

- *Clearing DHCP Bindings for Subscriber Access*
- *Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*
- clear dhcp server binding on page 198

List of Sample Output

Output Fields

Table 10 on page 191 lists the output fields for the **show dhcp server binding** command. Output fields are listed in the approximate order in which they appear.

Table 10: show dhcp server binding Output Fields

Field Name	Field Description	Level of Output
number clients, (*number* init, *number* bound, *number* selecting, *number* requesting, *number* renewing, *number* releasing)	Summary counts of the total number of DHCP clients and the number of DHCP clients in each state.	summary
IP address	IP address of the DHCP client.	brief detail
Session Id	Session ID of the subscriber session.	brief detail
Hardware address	Hardware address of the DHCP client.	brief detail
Expires	Number of seconds in which lease expires.	brief detail
State	State of the address binding table on the extended DHCP local server: • **BOUND**—Client has active IP address lease. • **FORCERENEW**—Client has received forcerenew message from server. • **INIT**—Initial state. • **RELEASE**—Client is releasing IP address lease. • **RENEWING**—Client sending request to renew IP address lease. • **REQUESTING**—Client requesting a DHCP server. • **SELECTING**—Client receiving offers from DHCP servers.	brief detail
Interface	Interface on which the request was received.	brief

Table 10: show dhcp server binding Output Fields *(continued)*

Field Name	Field Description	Level of Output
Lease Expires	Date and time at which the client's IP address lease expires.	**detail**
Lease Expires in	Number of seconds in which lease expires.	**detail**
Lease Start	Date and time at which the client's IP address lease started.	**detail**
Last Packet Received	Date and time at which the router received the last packet.	**detail**
Incoming Client Interface	Client's incoming interface.	**detail**
Client Interface Svlan Id	S-VLAN ID of the client's incoming interface.	**detail**
Client Interface Vlan Id	VLAN ID of the client's incoming interface.	**detail**
Demux Interface	Name of the IP demultiplexing (demux) interface.	**detail**
Server IP Address or **Server Identifier**	IP address of DHCP server.	**detail**
Server Interface	Interface of DHCP server.	**detail**
Client Pool Name	Name of address pool used to assign client IP address lease.	**detail**
ACI Interface Set Name	Internally generated name of the dynamic agent circuit identifier (ACI) interface set.	**detail**
ACI Interface Set Index	Index number of the dynamic ACI interface set.	**detail**
ACI Interface Set Session ID	Identifier of the dynamic ACI interface set entry in the session database.	**detail**

Sample Output

show dhcp server binding

```
user@host> show dhcp server binding
IP address        Session Id  Hardware address   Expires     State      Interface
100.20.20.15      6           00:10:94:00:00:01  86180       BOUND      ge-1/0/0.0

100.20.20.16      7           00:10:94:00:00:02  86180       BOUND      ge-1/0/0.0

100.20.20.17      8           00:10:94:00:00:03  86180       BOUND      ge-1/0/0.0

100.20.20.18      9           00:10:94:00:00:04  86180       BOUND      ge-1/0/0.0

100.20.20.19      10          00:10:94:00:00:05  86180       BOUND      ge-1/0/0.0
```

show dhcp server binding detail

```
user@host> show dhcp server binding detail
```

```
Client IP Address:  100.20.20.15
     Hardware Address:             00:10:94:00:00:01
     State:                       BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)

     Lease Expires:                2009-07-21 10:10:25 PDT
     Lease Expires in:             86151 seconds
     Lease Start:                  2009-07-20 10:10:25 PDT
     Incoming Client Interface:    ge-1/0/0.0
     Server Ip Address:            100.20.20.9
     Server Interface:             none
     Session Id:                   6
     Client Pool Name:             6
     Client IP Address:            100.20.20.16
     Hardware Address:             00:10:94:00:00:02
     State:                       BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)

     Lease Expires:                2009-07-21 10:10:25 PDT
     Lease Expires in:             86151 seconds
     Lease Start:                  2009-07-20 10:10:25 PDT
     Incoming Client Interface:    ge-1/0/0.0
     Server Ip Address:            100.20.20.9
     Server Interface:             none
     Session Id:                   7
     Client Pool Name:             7
```

show dhcp server binding detail (ACI Interface Set Configured)

```
user@host> show dhcp server binding detail
Client IP Address:  100.20.22.14
     Hardware Address:             00:00:64:34:01:02
     State:                        BOUND(LOCAL_SERVER_STATE_BOUND)
     Lease Expires:                2012-03-13 09:53:32 PDT
     Lease Expires in:             82660 seconds
     Lease Start:                  2012-03-12 10:23:32 PDT
     Last Packet Received:         2012-03-12 10:23:32 PDT
     Incoming Client Interface:    demux0.1073741827
     Client Interface Svlan Id:    1802
     Client Interface Vlan Id:     302
     Demux Interface:              demux0.1073741832
     Server Identifier:            100.20.200.202
     Session Id:                   11
     Client Pool Name:             poolA
     Client Profile Name:          DEMUXprofile
     ACI Interface Set Name:       aci-1002-demux0.1073741827
     ACI Interface Set Index:      2
     ACI Interface Set Session ID: 6
```

show dhcp server binding interface <vlan-id>

```
user@host> show dhcp server binding interface ge-1/1/0:100
IP address        Session Id  Hardware address   Expires     State      Interface
200.20.20.15       6           00:10:94:00:00:01  86124       BOUND
ge-1/1/0:100
```

show dhcp server binding interface <svlan-id>

```
user@host> show dhcp server binding interface ge-1/1/0:10-100
IP address        Session Id  Hardware address   Expires     State      Interface
200.20.20.16       7           00:10:94:00:00:02  86124       BOUND
ge-1/1/0:10-100
```

show dhcp server binding <ip-address>

```
user@host> show dhcp server binding 100.20.20.19
IP address        Session Id  Hardware address   Expires     State      Interface
100.20.20.19      10          00:10:94:00:00:05  86081       BOUND      ge-1/0/0.0
```

show dhcp server binding <session-id>

```
user@host> show dhcp server binding 6
IP address        Session Id  Hardware address   Expires     State      Interface
200.20.20.15      6           00:10:94:00:00:01  86124       BOUND      ge-1/0/0.0
```

show dhcp server binding summary

```
user@host> show dhcp server binding summary
3 clients, (2 init, 1 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)
```

show dhcp server binding <interfaces-vlan>

```
user@host> show dhcp server binding ge-1/0/0:100-200
IP address        Session Id  Hardware address   Expires     State      Interface
192.168.0.17       42          00:10:94:00:00:02  86346       BOUND
ge-1/0/0.1073741827
192.168.0.16       41          00:10:94:00:00:01  86346       BOUND
ge-1/0/0.1073741827
```

show dhcp server binding <interfaces-wildcard>

```
user@host> show dhcp server binding ge-1/3/*
IP address        Session Id  Hardware address   Expires     State      Interface
192.168.0.9        24          00:10:94:00:00:04  86361       BOUND
ge-1/3/0.110
192.168.0.8        23          00:10:94:00:00:03  86361       BOUND
ge-1/3/0.110
192.168.0.7        22          00:10:94:00:00:02  86361       BOUND
ge-1/3/0.110
```

show dhcp server statistics

Syntax

show dhcp server statistics
<logical-system *logical-system-name*>
<routing-instance *routing-instance-name*>

Release Information

Command introduced in Junos OS Release 9.0.

Description

Display extended Dynamic Host Configuration Protocol (DHCP) local server statistics.

Options

logical-system *logical-system-name*—(Optional) Display information about extended DHCP local server statistics on the specified logical system. If you do not specify a logical system, statistics are displayed for the default logical system.

routing-instance *routing-instance-name*—(Optional) Display information about extended DHCP local server statistics on the specified routing instance. If you do not specify a routing instance, statistics are displayed for the default routing instance.

Required Privilege Level

view

Related Documentation

- clear dhcp server statistics on page 201

List of Sample Output

show dhcp server statistics on page 196

Output Fields

Table 11 on page 196 lists the output fields for the **show dhcp server statistics** command. Output fields are listed in the approximate order in which they appear.

Table 11: show dhcp server statistics Output Fields

Field Name	Field Description
Packets dropped	Number of packets discarded by the extended DHCP local server because of errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears. • **Total**—Total number of packets discarded by the extended DHCP local server • **Authentication**—Number of packets discarded because they could not be authenticated • **Bad hardware address**—Number of packets discarded because an invalid hardware address was specified • **Bad opcode**—Number of packets discarded because an invalid operation code was specified • **Bad options**—Number of packets discarded because invalid options were specified • **Dynamic profile**—Number of packets discarded due to dynamic profile information • **Invalid server address**—Number of packets discarded because an invalid server address was specified • **No available addresses**—Number of packets discarded because there were no addresses available for assignment • **No interface match**—Number of packets discarded because they did not belong to a configured interface • **No routing instance match**—Number of packets discarded because they did not belong to a configured routing instance • **No valid local address**—Number of packets discarded because there was no valid local address • **Packet too short**—Number of packets discarded because they were too short • **Read error**—Number of packets discarded because of a system read error • **Send error**—Number of packets that the extended DHCP local server could not send
Messages received	Number of DHCP messages received. • **BOOTREQUEST**—Number of BOOTP protocol data units (PDUs) received • **DHCPDECLINE**—Number of DHCP PDUs of type DECLINE received • **DHCPDISCOVER**—Number of DHCP PDUs of type DISCOVER received • **DHCPINFORM**—Number of DHCP PDUs of type INFORM received • **DHCPRELEASE**—Number of DHCP PDUs of type RELEASE received • **DHCPREQUEST**—Number of DHCP PDUs of type REQUEST received
Messages sent	Number of DHCP messages sent. • **BOOTREPLY**—Number of BOOTP PDUs transmitted • **DHCPOFFER**—Number of DHCP OFFER PDUs transmitted • **DHCPACK**—Number of DHCP ACK PDUs transmitted • **DHCPNACK**—Number of DHCP NACK PDUs transmitted • **DHCPFORCERENEW**—Number of DHCP FORCERENEW PDUs transmitted

Sample Output

show dhcp server statistics

```
user@host> show dhcp server statistics
Packets dropped:
    Total                     0

Messages received:
```

```
    BOOTREQUEST                 25
    DHCPDECLINE                 0
    DHCPDISCOVER                10
    DHCPINFORM                  0
    DHCPRELEASE                 4
    DHCPREQUEST                 10

Messages sent:
    BOOTREPLY                   20
    DHCPOFFER                   10
    DHCPACK                     10
    DHCPNAK                     0
    DHCPFORCERENEW              0
```

clear dhcp server binding

Syntax

clear dhcp server binding
<address>
<all>
<interface *interface-name*>
<interfaces-vlan>
<interfaces-wildcard>
<logical-system *logical-system-name*>
<routing-instance *routing-instance-name*>

Release Information

Command introduced in Junos OS Release 9.0.
Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

Description

Clear the binding state of a Dynamic Host Configuration Protocol (DHCP) client from the client table on the extended DHCP local server.

Options

address—(Optional) Clear the binding state for the DHCP client, using one of the following entries:

- *ip-address*—The specified IP address.
- *mac-address*—The specified MAC address.
- *session-id*—The specified session ID.

all—(Optional) Clear the binding state for all DHCP clients.

interface *interface-name*—(Optional) Clear the binding state for DHCP clients on the specified interface.

NOTE: This option clears all bindings whose initial login requests were received over the specified interface. Dynamic demux login requests are not received over the dynamic demux interface, but rather the underlying interface of the dynamic demux interface. To clear a specific dynamic demux interface, use the ***ip-address*** or ***mac-address*** options.

interfaces-vlan—(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.

interfaces-wildcard—(Optional) Clear bindings on a set of interfaces. This option supports the use of the wildcard character (*).

logical-system *logical-system-name*—(Optional) Clear the binding state for DHCP clients on the specified logical system.

routing-instance *routing-instance-name*—(Optional) Clear the binding state for DHCP clients on the specified routing instance.

Required Privilege Level view

Related Documentation
- *Clearing DHCP Bindings for Subscriber Access*
- show dhcp server binding on page 190

List of Sample Output

Output Fields See show dhcp server binding for an explanation of output fields.

Sample Output

clear dhcp server binding <ip-address>

The following sample output displays the address bindings in the DHCP client table on the extended DHCP local server before and after the **clear dhcp server binding** command is issued.

```
user@host> show dhcp server binding

2 clients, (0 bound, 0 selecting, 0 renewing, 0 rebinding)

IP address       Hardware address   Type     Lease expires at
100.20.32.1      90:00:00:01:00:01  active   2007-01-17 11:38:47 PST
100.20.32.3      90:00:00:02:00:01  active   2007-01-17 11:38:41 PST


user@host> clear dhcp server binding 10.20.32.1

user@host> show dhcp server binding

1 clients, (0 bound, 0 selecting, 0 renewing, 0 rebinding)

IP address       Hardware address   Type     Lease expires at
100.20.32.3      90:00:00:02:00:01  active   2007-01-17 11:38:41 PST
```

clear dhcp server binding all

The following command clears all DHCP local server bindings:

```
user@host> clear dhcp server binding all
```

clear dhcp server binding interface

The following command clears DHCP local server bindings on a specific interface:

```
user@host> clear dhcp server binding interface fe-0/0/2
```

clear dhcp server binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCP local server bindings on top of the underlying interface **ae0**, which clears DHCP bindings on all demux VLANs on top of **ae0**:

```
user@host> clear dhcp server binding ae0
```

clear dhcp server binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCP local server bindings over a specific interface:

```
user@host> clear dhcp server binding ge-1/0/0.*
```

clear dhcp server statistics

Syntax

clear dhcp server statistics
<interface *interface-name*>
<logical-system *logical-system-name*>
<routing-instance *routing-instance-name*>

Release Information

Command introduced in Junos OS Release 9.0.

Description

Clear all extended Dynamic Host Configuration Protocol (DHCP) local server statistics.

Options

logical-system *logical-system-name*—(Optional) Clear the statistics for DHCP clients on the specified logical system. If you do not specify a logical system, statistics are cleared for the default logical system.

routing-instance *routing-instance-name*—(Optional) Clear the statistics for DHCP clients on the specified routing instance. If you do not specify a routing instance, statistics are cleared for the default routing instance.

Required Privilege Level

view

List of Sample Output

clear dhcp server statistics on page 201

Output Fields

See show dhcp server statistics for an explanation of output fields.

Sample Output

clear dhcp server statistics

The following sample output displays the extended DHCP local server statistics before and after the **clear dhcp server statistics** command is issued.

```
user@host> show dhcp server statistics
Packets dropped:
    Total                       0

Messages received:
    BOOTREQUEST                 89163
    DHCPDECLINE                 0
    DHCPDISCOVER                8110
    DHCPINFORM                  0
    DHCPRELEASE                 0
    DHCPREQUEST                 81053

Messages sent:
    BOOTREPLY                   32420
    DHCPOFFER                   8110
    DHCPACK                     8110
    DHCPNAK                     8100


user@host> clear dhcp server statistics

user@host> show dhcp server statistics
```

```
Packets dropped:
    Total                       0

Messages received:
    BOOTREQUEST                 0
    DHCPDECLINE                 0
    DHCPDISCOVER                0
    DHCPINFORM                  0
    DHCPRELEASE                 0
    DHCPREQUEST                 0

Messages sent:
    BOOTREPLY                   0
    DHCPOFFER                   0
    DHCPACK                     0
    DHCPNAK                     0
```

CHAPTER 9

Subscriber Management DHCP Relay CLI Commands

show dhcp relay binding

Syntax

show dhcp relay binding
<*address*>
<**brief**>
<**detail**>
<**interface** *interface-name*>
<*interfaces-vlan*>
<*interfaces-wildcard*>
<*ip-address* | *mac-address*>
<**logical-system** *logical-system-name*>
<**routing-instance** *routing-instance-name*>
<**summary**>

Release Information

Command introduced in Junos OS Release 8.3.
Options **interface** and ***mac-address*** added in Junos OS Release 8.4.
Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.
Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.

Description

Display the address bindings in the Dynamic Host Configuration Protocol (DHCP) client table.

Options

address—(Optional) Display DHCP binding information for a specific client identified by one of the following entries:

- *ip-address*—The specified IP address.
- *mac-address*—The specified MAC address.
- *session-id*—The specified session ID.

brief—(Optional) Display brief information about the active client bindings. This is the default, and produces the same output as **show dhcp relay binding**.

detail—(Optional) Display detailed client binding information.

interface ***interface-name***—(Optional) Perform this operation on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.

interfaces-vlan—(Optional) Show the binding state information on the interface VLAN ID and S-VLAN ID.

interfaces-wildcard—(Optional) The set of interfaces on which to show binding state information. This option supports the use of the wildcard character (*).

logical-system ***logical-system-name***—(Optional) Perform this operation on the specified logical system.

routing-instance ***routing-instance-name***—(Optional) Perform this operation on the specified routing instance.

summary—(Optional) Display a summary of DHCP client information.

Required Privilege Level view

Related Documentation
- *Clearing DHCP Bindings for Subscriber Access*
- clear dhcp relay binding on page 212

List of Sample Output
show dhcp relay binding on page 206
show dhcp relay binding detail on page 206
show dhcp relay binding interface on page 207
show dhcp relay binding interface vlan-id on page 207
show dhcp relay binding interface svlan-id on page 207
show dhcp relay binding ip-address on page 207
show dhcp relay binding mac-address on page 207
show dhcp relay binding session-id on page 207
show dhcp relay binding <interfaces-vlan> on page 208
show dhcp relay binding <interfaces-wildcard> on page 208
show dhcp relay binding summary on page 208

Output Fields Table 12 on page 205 lists the output fields for the **show dhcp relay binding** command. Output fields are listed in the approximate order in which they appear.

Table 12: show dhcp relay binding Output Fields

Field Name	Field Description	Level of Output
number clients,(*number* init, *number* bound, *number* selecting, *number* requesting, *number* renewing,*number* rebinding, *number* releasing)	Summary counts of the total number of DHCP clients and the number of DHCP clients in each state.	summary
IP address	IP address of the DHCP client.	briefdetail
Session Id	Session ID of the subscriber session.	briefdetail
Hardware address	Hardware address of the DHCP client.	briefdetail
Expires	Number of seconds in which the lease expires.	briefdetail
State	State of the DHCP relay address binding table on the DHCP client: • **BOUND**—Client has an active IP address lease. • **INIT**—Initial state. • **REBINDING**—Client is broadcasting a request to renew the IP address lease. • **RELEASE**—Client is releasing the IP address lease. • **RENEWING**—Client is sending a request to renew the IP address lease. • **REQUESTING**—Client is requesting a DHCP server. • **SELECTING**—Client is receiving offers from DHCP servers.	briefdetail

Table 12: show dhcp relay binding Output Fields *(continued)*

Field Name	Field Description	Level of Output
Interface	Incoming client interface.	brief
Lease Expires	Date and time at which the client's IP address lease expires.	detail
Lease Expires in	Number of seconds in which the lease expires.	detail
Lease Start	Date and time at which the client's IP address lease started.	detail
Incoming Client Interface	Client's incoming interface.	detail
Server IP Address	IP address of the DHCP server.	detail
Server Interface	Interface of the DHCP server.	detail
Bootp Relay Address	IP address of BOOTP relay.	detail
Type	Type of DHCP packet processing performed on the router: • **active**—Router actively processes and relays DHCP packets. • **passive**—Router passively snoops DHCP packets passing through the router.	All levels
Lease expires at	Date and time at which the client's IP address lease expires.	All levels

Sample Output

show dhcp relay binding

```
user@host> show dhcp relay binding
IP address        Session Id  Hardware address   Expires     State      Interface
100.20.32.11      41          00:10:94:00:00:01  86371       BOUND      ge-1/0/0.0

100.20.32.12      42          00:10:94:00:00:02  86371       BOUND      ge-1/0/0.0

100.20.32.13      43          00:10:94:00:00:03  86371       BOUND      ge-1/0/0.0

100.20.32.14      44          00:10:94:00:00:04  86371       BOUND      ge-1/0/0.0

100.20.32.15      45          00:10:94:00:00:05  86371       BOUND      ge-1/0/0.0
```

show dhcp relay binding detail

```
user@host> show dhcp relay binding detail


Client IP Address:  100.20.32.11
     Hardware Address:             00:10:94:00:00:01
     State:                        BOUND(DHCP_RELAY_STATE_BOUND_ON_INTF_DELETE)
     Lease Expires:                2009-07-21 11:00:06 PDT
     Lease Expires in:             86361 seconds
```

```
    Lease Start:                  2009-07-20 11:00:06 PDT
    Last Packet Received:         2009-07-20 11:00:06 PDT
    Incoming Client Interface:    ge-1/0/0.0
    Server Ip Address:            100.20.22.2
    Server Interface:             none
    Bootp Relay Address:          100.20.32.2
    Session Id:                   41

Client IP Address:  100.20.32.12
    Hardware Address:             00:10:94:00:00:02
    State:                        BOUND(DHCP_RELAY_STATE_BOUND_ON_INTF_DELETE)
    Lease Expires:                2009-07-21 11:00:06 PDT
    Lease Expires in:             86361 seconds
    Lease Start:                  2009-07-20 11:00:06 PDT
    Last Packet Received:         2009-07-20 11:00:06 PDT
    Incoming Client Interface:    ge-1/0/0.0
    Server Ip Address:            100.20.22.2
    Server Interface:             none
    Bootp Relay Address:          100.20.32.2
    Session Id:                   42
```

show dhcp relay binding interface

```
user@host> show dhcp relay binding interface fe-0/0/2

IP address       Hardware address   Type     Lease expires at
100.20.32.1      90:00:00:01:00:01  active   2007-03-27 15:06:20 EDT
```

show dhcp relay binding interface vlan-id

```
user@host> show dhcp relay binding interface ge-1/1/0:100

IP address        Session Id  Hardware address   Expires     State      Interface
200.20.20.15       6           00:10:94:00:00:01  86124       BOUND
ge-1/1/0:100
```

show dhcp relay binding interface svlan-id

```
user@host> show dhcp relay binding interface ge-1/1/0:10-100

IP address        Session Id  Hardware address   Expires     State      Interface
200.20.20.16       7           00:10:94:00:00:02  86124       BOUND
ge-1/1/0:10-100
```

show dhcp relay binding ip-address

```
user@host> show dhcp relay binding 100.20.32.13
IP address        Session Id  Hardware address   Expires     State      Interface
100.20.32.13      43         00:10:94:00:00:03  86293       BOUND      ge-1/0/0.0
```

show dhcp relay binding mac-address

```
user@host> show dhcp relay binding 00:10:94:00:00:05
IP address        Session Id  Hardware address   Expires     State      Interface
100.20.32.15      45         00:10:94:00:00:05  86279       BOUND      ge-1/0/0.0
```

show dhcp relay binding session-id

```
user@host> show dhcp relay binding 41
```

```
 IP address        Session Id  Hardware address   Expires     State      Interface
100.20.32.11      41          00:10:94:00:00:01  86305       BOUND      ge-1/0/0.0
```

show dhcp relay binding <interfaces-vlan>

```
user@host> show dhcp relay binding ge-1/0/0:100-200
IP address        Session Id  Hardware address   Expires     State      Interface
192.168.0.17       42          00:10:94:00:00:02  86346       BOUND
ge-1/0/0.1073741827
192.168.0.16       41          00:10:94:00:00:01  86346       BOUND
ge-1/0/0.1073741827
```

show dhcp relay binding <interfaces-wildcard>

```
user@host> show dhcp relay binding ge-1/3/*
IP address        Session Id  Hardware address   Expires     State      Interface
192.168.0.9        24          00:10:94:00:00:04  86361       BOUND
ge-1/3/0.110
192.168.0.8        23          00:10:94:00:00:03  86361       BOUND
ge-1/3/0.110
192.168.0.7        22          00:10:94:00:00:02  86361       BOUND
ge-1/3/0.110
```

show dhcp relay binding summary

```
user@host> show dhcp relay binding summary
3 clients, (2 init, 1 bound, 0 selecting, 0 requesting, 0 renewing, 0 rebinding,
 0 releasing)
```

show dhcp relay statistics

Syntax

show dhcp relay statistics
<logical-system *logical-system-name*>
<routing-instance *routing-instance-name*>

Syntax

Syntax for EX Series switches:

show dhcp relay statistics
<routing-instance *routing-instance-name*>

Release Information

Command introduced in Junos OS Release 8.3.
Command introduced in Junos OS Release 12.1 for EX Series switches.
Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.

Description

Display Dynamic Host Configuration Protocol (DHCP) relay statistics.

Options

logical-system *logical-system-name*—(On routers only) (Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are displayed for the default logical system.

routing-instance *routing-instance-name*—(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are displayed for the default routing instance.

Required Privilege Level

view

Related Documentation

- clear dhcp relay statistics on page 214

List of Sample Output

show dhcp relay statistics on page 211

Output Fields

Table 13 on page 210 lists the output fields for the **show dhcp relay statistics** command. Output fields are listed in the approximate order in which they appear.

Table 13: show dhcp relay statistics Output Fields

Field Name	Field Description
Packets dropped	Number of packets discarded by the extended DHCP relay agent application due to errors. Only nonzero statistics appear in the **Packets dropped** output. When all of the Packets dropped statistics are 0 (zero), only the **Total** field appears. • **Total**—Total number of packets discarded by the extended DHCP relay agent application. • **Bad hardware address**—Number of packets discarded because an invalid hardware address was specified. • **Bad opcode**—Number of packets discarded because an invalid operation code was specified. • **Bad options**—Number of packets discarded because invalid options were specified. • **Invalid server address**—Number of packets discarded because an invalid server address was specified. • **No available addresses**—Number of packets discarded because there were no addresses available for assignment. • **No interface match**—Number of packets discarded because they did not belong to a configured interface. • **No routing instance match**—Number of packets discarded because they did not belong to a configured routing instance. • **No valid local address**—Number of packets discarded because there was no valid local address. • **Packet too short**—Number of packets discarded because they were too short. • **Read error**—Number of packets discarded because of a system read error. • **Send error**—Number of packets that the extended DHCP relay application could not send. • **Option 60**—Number of packets discarded containing DHCP option 60 vendor-specific information. • **Option 82**—Number of packets discarded because DHCP option 82 information could not be added.
Messages received	Number of DHCP messages received. • **BOOTREQUEST**—Number of BOOTP protocol data units (PDUs) received • **DHCPDECLINE**—Number of DHCP PDUs of type DECLINE received • **DHCPDISCOVER**—Number of DHCP PDUs of type DISCOVER received • **DHCPINFORM**—Number of DHCP PDUs of type INFORM received • **DHCPRELEASE**—Number of DHCP PDUs of type RELEASE received • **DHCPREQUEST**—Number of DHCP PDUs of type REQUEST received
Messages sent	Number of DHCP messages sent. • **BOOTREPLY**—Number of BOOTP PDUs transmitted • **DHCPOFFER**—Number of DHCP OFFER PDUs transmitted • **DHCPACK**—Number of DHCP ACK PDUs transmitted • **DHCPNACK**—Number of DHCP NACK PDUs transmitted • **DHCPFORCERENEW**—Number of DHCP FORCERENEW PDUs transmitted
Packets forwarded	Number of packets forwarded. • **BOOTREQUEST**—Number of BOOTREQUEST protocol data units (PDUs) forwarded • **BOOTREPLY**—Number of BOOTREPLY protocol data units (PDUs) forwarded

Sample Output

show dhcp relay statistics

```
user@host> show dhcp relay statistics
Packets dropped:
    Total                       30
    Bad hardware address        1
    Bad opcode                  1
    Bad options                 3
    Invalid server address      5
    No available addresses      1
    No interface match          2
    No routing instance match   9
    No valid local address      4
    Packet too short            2
    Read error                  1
    Send error                  1
    Option 60                   1
    Option 82                   2

Messages received:
    BOOTREQUEST                 116
    DHCPDECLINE                 0
    DHCPDISCOVER                11
    DHCPINFORM                  0
    DHCPRELEASE                 0
    DHCPREQUEST                 105

Messages sent:
    BOOTREPLY                   0
    DHCPOFFER                   2
    DHCPACK                     1
    DHCPNAK                     0
    DHCPFORCERENEW              0

Packets forwarded:
    Total                       4
    BOOTREQUEST                 2
    BOOTREPLY                   2
```

clear dhcp relay binding

Syntax

clear dhcp relay binding
<*address*>
<**all**>
<**interface** *interface-name*>
<*interfaces-vlan*>
<*interfaces-wildcard*>
<**logical-system** *logical-system-name*>
<**routing-instance** *routing-instance-name*>

Release Information

Command introduced in Junos OS Release 8.3.
Options **all** and **interface** added in Junos OS Release 8.4.
Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.
Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.

Description

Clear the binding state of a Dynamic Host Configuration Protocol (DHCP) client from the client table.

Options

address—(Optional) Clear the binding state for the DHCP client, using one of the following entries:

- *ip-address*—The specified IP address.
- *mac-address*—The specified MAC address.
- *session-id*—The specified session ID.

all—(Optional) Clear the binding state for all DHCP clients.

interface ***interface-name***—(Optional) Clear the binding state for DHCP clients on the specified interface.

interfaces-vlan—(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.

interfaces-wildcard—(Optional) The set of interfaces on which to clear bindings. This option supports the use of the wildcard character (*).

logical-system ***logical-system-name***—(Optional) Clear the binding state for DHCP clients on the specified logical system.

routing-instance ***routing-instance-name***—(Optional) Clear the binding state for DHCP clients on the specified routing instance.

Required Privilege Level

view

Related Documentation

- *Clearing DHCP Bindings for Subscriber Access*
- show dhcp relay binding on page 204

List of Sample Output

Output Fields See show dhcp relay binding for an explanation of output fields.

Sample Output

clear dhcp relay binding

The following sample output displays the address bindings in the DHCP client table before and after the **clear dhcp relay binding** command is issued.

```
user@host> show dhcp relay binding
IP address       Hardware address   Type     Lease expires at
100.20.32.1      90:00:00:01:00:01  active   2007-02-08 16:41:17 EST
192.168.14.8      90:00:01:01:02:01  active   2007-02-10 10:01:06 EST


user@host> clear dhcp relay binding 100.20.32.1


user@host> show dhcp relay binding
IP address       Hardware address   Type     Lease expires at
192.168.14.8      90:00:01:01:02:01  active   2007-02-10 10:01:06 EST
```

clear dhcp relay binding all

The following command clears all DHCP relay agent bindings:

```
user@host> clear dhcp relay binding all
```

clear dhcp relay binding interface

The following command clears DHCP relay agent bindings on a specific interface:

```
user@host> clear dhcp relay binding interface fe-0/0/3
```

clear dhcp relay binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCP relay agent bindings on top of the underlying interface **ae0**, which clears DHCP bindings on all demux VLANs on top of **ae0**:

```
user@host> clear dhcp relay binding interface ae0
```

clear dhcp relay binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCP relay agent bindings over a specific interface:

```
user@host> clear dhcp relay binding ge-1/0/0.*
```

clear dhcp relay statistics

Syntax

clear dhcp relay statistics
<logical-system *logical-system-name*>
<routing-instance *routing-instance-name*>

Syntax

Syntax for EX Series switches:

show dhcp relay statistics
<routing-instance *routing-instance-name*>

Release Information

Command introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.

Description

Clear all Dynamic Host Configuration Protocol (DHCP) relay statistics.

Options

logical-system *logical-system-name*—(On routers only) (Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are cleared for the default logical system.

routing-instance *routing-instance-name*—(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are cleared for the default routing instance.

Required Privilege Level

view

Related Documentation

- show dhcp relay statistics on page 209

List of Sample Output

clear dhcp relay statistics on page 215

Output Fields

Table 14 on page 215 lists the output fields for the **clear dhcp relay statistics** command.

Table 14: clear dhcp relay statistics Output Fields

Field Name	Field Description
Packets dropped	Number of packets discarded by the extended DHCP relay agent application due to errors. Only nonzero statistics appear in the **Packets dropped** output. When all of the Packets dropped statistics are 0 (zero), only the **Total** field appears. • **Total**—Total number of packets discarded by the extended DHCP relay agent application. • **Bad hardware address**—Number of packets discarded because an invalid hardware address was specified. • **Bad opcode**—Number of packets discarded because an invalid operation code was specified. • **Bad options**—Number of packets discarded because invalid options were specified. • **Invalid server address**—Number of packets discarded because an invalid server address was specified. • **No available addresses**—Number of packets discarded because there were no addresses available for assignment. • **No interface match**—Number of packets discarded because they did not belong to a configured interface. • **No routing instance match**—Number of packets discarded because they did not belong to a configured routing instance. • **No valid local address**—Number of packets discarded because there was no valid local address. • **Packet too short**—Number of packets discarded because they were too short. • **Read error**—Number of packets discarded because of a system read error. • **Send error**—Number of packets that the extended DHCP relay application could not send. • **Option 60**—Number of packets discarded containing DHCP option 60 vendor-specific information. • **Option 82**—Number of packets discarded because DHCP option 82 information could not be added.
Messages received	Number of DHCP messages received. • **BOOTREQUEST**—Number of BOOTP protocol data units (PDUs) received • **DHCPDECLINE**—Number of DHCP PDUs of type DECLINE received • **DHCPDISCOVER**—Number of DHCP PDUs of type DISCOVER received • **DHCPINFORM**—Number of DHCP PDUs of type INFORM received • **DHCPRELEASE**—Number of DHCP PDUs of type RELEASE received • **DHCPREQUEST**—Number of DHCP PDUs of type REQUEST received
Messages sent	Number of DHCP messages sent. • **BOOTREPLY**—Number of BOOTP PDUs transmitted • **DHCPOFFER**—Number of DHCP OFFER PDUs transmitted • **DHCPACK**—Number of DHCP ACK PDUs transmitted • **DHCPNACK**—Number of DHCP NACK PDUs transmitted

Sample Output

clear dhcp relay statistics

The following sample output displays the DHCP relay statistics before and after the **clear dhcp relay statistics** command is issued.

```
user@host> show dhcp relay statistics
```

```
Packets dropped:
    Total                       0

Messages received:
    BOOTREQUEST                 116
    DHCPDECLINE                 0
    DHCPDISCOVER                11
    DHCPINFORM                  0
    DHCPRELEASE                 0
    DHCPREQUEST                 105

Messages sent:
    BOOTREPLY                   44
    DHCPOFFER                   11
    DHCPACK                     11
    DHCPNAK                     11
```

```
user@host> clear dhcp relay statistics
```

```
user@host> show dhcp relay statistics
Packets dropped:
    Total                       0

Messages received:
    BOOTREQUEST                 0
    DHCPDECLINE                 0
    DHCPDISCOVER                0
    DHCPINFORM                  0
    DHCPRELEASE                 0
    DHCPREQUEST                 0

Messages sent:
    BOOTREPLY                   0
    DHCPOFFER                   0
    DHCPACK                     0
    DHCPNAK                     0
```

CHAPTER 10

Subscriber Management Interface CLI Commands

show interfaces (Loopback)

Syntax

show interfaces lo0
<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 local loopback interface.

NOTE: Logical interface lo0.16385 is the loopback interface for the internal routing instance. Created by the internal routing service process, this interface facilitates internal traffic. It prevents any filter created on loopback lo0.0 from blocking internal traffic.

Options

lo0—Display standard status information about the local loopback interface.

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

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

List of Sample Output

show interfaces (Loopback) on page 221
show interfaces brief (Loopback) on page 222
show interfaces detail (Loopback) on page 222
show interfaces extensive (Loopback) on page 223

Output Fields

Table 15 on page 218 lists the output fields for the **show interfaces** (loopback) command. Output fields are listed in the approximate order in which they appear.

Table 15: Loopback show interfaces Output Fields

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

Table 15: Loopback show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Enabled	State of the interface. Possible values are described in the "Enabled Field" section under *Common Output Fields Description*.	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
Link-level type	Encapsulation type used on the physical interface.	All levels
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*.	All levels
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under *Common Output Fields Description*.	All levels
Link type	Data transmission type.	**detail extensive**
Link flags	Information about the link. Possible values are described in the "Link Flags" section under *Common Output Fields Description*.	**detail extensive** none
Physical info	Information about the physical interface.	**detail extensive**
Hold-times	Current interface hold-time up and hold-time down. Value is in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive**
Hardware address	Media access control (MAC) address of the interface.	**detail extensive**
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**
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**

Table 15: Loopback 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	• **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 the frames 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. • **MTU errors**—Number of packets larger than the MTU threshold. • **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**
SNMP ifIndex	Logical interface SNMP interface index number.	**detail extensive**
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*.	**brief detail extensive**
Encapsulation	Encapsulation on the logical interface.	**brief detail extensive**
Input packets	Number of packets received on the logical interface.	None specified

Table 15: Loopback 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	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**
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**
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*.	**detail extensive** none
Addresses, Flags	Information about the address flags. Possible values are described in the "Addresses Flags" section under *Common Output Fields Description*.	**detail extensive**
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 on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces (Loopback)

```
user@host> show interfaces lo0
Physical interface: lo0, Enabled, Physical link is Up
  Interface index: 6, SNMP ifIndex: 6
  Type: Loopback, MTU: Unlimited
  Device flags   : Present Running Loopback
  Interface flags: SNMP-Traps
  Link flags     : None
  Last flapped   : Never
    Input packets : 0
    Output packets: 0
```

```
  Logical interface lo0.0 (Index 64) (SNMP ifIndex 16)
    Flags: SNMP-Traps Encapsulation: Unspecified
    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: Unlimited
      Flags: None
      Addresses, Flags: Is-Default Is-Primary
        Local: 10.0.0.1
      Addresses
        Local: 127.0.0.1
    Protocol iso, MTU: Unlimited
      Flags: None
      Addresses, Flags: Is-Default Is-Primary
        Local: 49.0004.1000.0000.0001

  Logical interface lo0.16385 (Index 65) (SNMP ifIndex 76)
    Flags: SNMP-Traps Encapsulation: Unspecified
    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: Unlimited
      Flags: None
```

show interfaces brief (Loopback)

```
user@host> show interfaces lo0 brief
Physical interface: lo0, Enabled, Physical link is Up
  Type: Loopback, Link-level type: Unspecified, MTU: Unlimited,
  Clocking: Unspecified, Speed: Unspecified
  Device flags   : Present Running Loopback
  Interface flags: SNMP-Traps

  Logical interface lo0.0
    Flags: SNMP-Traps Encapsulation: Unspecified
    inet  10.0.0.1         --> 0/0
          127.0.0.1        --> 0/0
    iso   49.0004.1000.0000.0001

  Logical interface lo0.16385
    Flags: SNMP-Traps Encapsulation: Unspecified
    inet
```

show interfaces detail (Loopback)

```
user@host> show interfaces lo0 detail
Physical interface: lo0, Enabled, Physical link is Up
  Interface index: 6, SNMP ifIndex: 6, Generation: 4
  Type: Loopback, Link-level type: Unspecified, MTU: Unlimited,
  Clocking: Unspecified, Speed: Unspecified
  Device flags   : Present Running Loopback
  Interface flags: 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 lo0.0 (Index 64) (SNMP ifIndex 16) (Generation 3)
   Flags: SNMP-Traps Encapsulation: Unspecified
   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

   Protocol inet, MTU: Unlimited, Generation: 10, Route table: 0
     Flags: None
     Addresses, Flags: Is-Default Is-Primary
       Destination: Unspecified, Local: 10.0.0.1, Broadcast: Unspecified,
       Generation: 10
     Addresses, Flags: None
       Destination: Unspecified, Local: 127.0.0.1, Broadcast: Unspecified,
       Generation: 12
   Protocol iso, MTU: Unlimited, Generation: 11, Route table: 0
     Flags: None
     Addresses, Flags: Is-Default Is-Primary
       Destination: Unspecified, Local: 49.0004.1000.0000.0001,
       Broadcast: Unspecified, Generation: 14

 Logical interface lo0.16385 (Index 65) (SNMP ifIndex 76) (Generation 4)
   Flags: SNMP-Traps Encapsulation: Unspecified
   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
   Protocol inet, MTU: Unlimited, Generation: 12, Route table: 1
     Flags: None
```

show interfaces extensive (Loopback)

```
user@host> show interfaces lo0 extensive
Physical interface: lo0, Enabled, Physical link is Up
  Interface index: 6, SNMP ifIndex: 6, Generation: 4
  Type: Loopback, Link-level type: Unspecified, MTU: Unlimited,
  Clocking: Unspecified, Speed: Unspecified
  Device flags   : Present Running Loopback
  Interface flags: 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 lo0.0 (Index 64) (SNMP ifIndex 16) (Generation 3)
    Flags: SNMP-Traps Encapsulation: Unspecified
    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
    Protocol inet, MTU: Unlimited, Generation: 10, Route table: 0
      Flags: None
      Addresses, Flags: Is-Default Is-Primary
        Destination: Unspecified, Local: 10.0.0.1, Broadcast: Unspecified,
        Generation: 10
      Addresses, Flags: None
        Destination: Unspecified, Local: 127.0.0.1, Broadcast: Unspecified,
        Generation: 12
    Protocol iso, MTU: Unlimited, Generation: 11, Route table: 0
      Flags: None
      Addresses, Flags: Is-Default Is-Primary
        Destination: Unspecified, Local: 49.0004.1000.0000.0001,
        Broadcast: Unspecified, Generation: 14

  Logical interface lo0.16385 (Index 65) (SNMP ifIndex 76) (Generation 4)
    Flags: SNMP-Traps Encapsulation: Unspecified
    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
    Protocol inet, MTU: Unlimited, Generation: 12, Route table: 1
      Flags: None
```

show interfaces (Aggregated Ethernet)

Syntax

show interfaces ae *number*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

(M Series, T Series, and MX Series routers and EX Series switches) Display status information about the specified aggregated Fast Ethernet or Gigabit Ethernet interface.

Options

ae *number*—Display standard information about the specified aggregated Fast Ethernet or Gigabit Ethernet interface.

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

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information.

snmp-index *snmp-index*—(Optional) Display information about the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

List of Sample Output

show interfaces (Aggregated Ethernet) on page 229
show interfaces brief (Aggregated Ethernet) on page 230
show interfaces detail (Aggregated Ethernet) on page 230
show interfaces extensive (Aggregated Ethernet) on page 231
show interfaces extensive (Aggregated Ethernet with VLAN Stacking) on page 232

Output Fields

Table 16 on page 225 lists the output fields for the **show interfaces** (Aggregated Ethernet) command. Output fields are listed in the approximate order in which they appear.

Table 16: show interfaces (Aggregated Ethernet) Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface and state of the interface.	All levels
Enabled	State of the physical interface. Possible values are described in the "Enabled Field" section under *Common Output Fields Description*.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	All levels

Table 16: show interfaces (Aggregated Ethernet) 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**
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Loopback	Loopback status: **Enabled** or **Disabled**. If loopback is enabled, type of loopback: **Local** or **Remote**.	All levels
Source filtering	Source filtering status: **Enabled** or **Disabled**.	All levels
Flow control	Flow control status: **Enabled** or **Disabled**.	All levels
Minimum links needed	Number of child links that must be operational for the aggregate interface to be operational.	All levels
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under *Common Output Fields Description*.	All levels
Interface flags	Information about the interface. Possible values are described in the "Interfaces Flags" section under *Common Output Fields Description*.	All levels
Current address	Configured MAC address.	**detail extensive**
Hardware address	Hardware MAC address.	**detail extensive**
Last flapped	Date, time, and how long ago the interface went from down to up or from up to down. The format is **Last flapped:** ***year-month-day hours:minutes:seconds timezone*** **(*hours:minutes:seconds* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).**	**detail extensive**
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**

Table 16: show interfaces (Aggregated Ethernet) Output Fields *(continued)*

Field Name	Field Description	Level of Output
Traffic statistics	Number of and rate at which bytes and packets are received and transmitted on the physical interface. • **Input bytes**—Number of bytes and rate, in bps, at which bytes are received on the interface. • **Output bytes**—Number of bytes and rate, in bps, at which bytes are transmitted on the interface. • **Input packets**—Number of packets and rate, in pps, at which packets are received on the interface. • **Output packets**—Number of packets and rate, in pps, at which packets are transmitted on the interface.	**detail extensive**
Input errors	Input errors on the interface: • **Errors**—Sum of incoming frame aborts and frame check sequence (FCS) errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's random early detection (RED) mechanism. • **Framing errors**—Number of packets received with an invalid FCS. • **Runts**—Number of frames received that are smaller than the runt threshold. • **Giants**—Number of frames received that are larger than the giant threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or were not of interest. Usually, this field reports protocols that Junos OS does not handle. • **Resource errors**—Sum of transmit drops.	**detail extensive**
Output errors	Output errors on the interface: • **Carrier transitions** —Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), then the cable, the far-end system, or the PIC is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**detail extensive**
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**

Table 16: show interfaces (Aggregated Ethernet) Output Fields *(continued)*

Field Name	Field Description	Level of Output
Queue counters	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface (which reflects its initialization sequence).	**detail extensive** none
SNMP ifIndex	SNMP interface index number of the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags Field" section under *Common Output Fields Description*.	All levels
VLAN-Tag	Tag Protocol Identifier (TPID) and VLAN identifier.	All levels
Demux	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: • **Source Family Inet** • **Destination Family Inet**	**detail extensive** none
Encapsulation	Encapsulation on the logical interface.	All levels

Statistics	Information about the number of packets, packets per second, number of bytes, and bytes per second on this aggregate interface.
	• **Bundle**—Information about input and output bundle rates.
	• **Link**—(**detail** and **extensive** only) Information about specific links in the aggregate, including link state and input and output rates.
	• **Marker Statistics**—(**detail** and **extensive** only) Information about 802.3ad marker protocol statistics on the specified links.
	  • **Marker Rx**—Number of valid marker protocol data units (PDUs) received on this aggregation port.
	  • **Resp Tx**—Number of marker response PDUs transmitted on this aggregation port.
	  • **Unknown Rx**—Number of frames received that either carry the slow protocols Ethernet type value (43B.4) but contain an unknown PDU, or are addressed to the slow protocols group MAC address (43B.3) but do not carry the slow protocols Ethernet type.
	  • **Illegal Rx**—Number of frames received that carry the slow protocols Ethernet type value (43B.4) but contain a badly formed PDU or an illegal value of protocol subtype (43B.4).
		detail extensive none

protocol-family	Protocol family configured on the logical interface. Possible values are described in the "Protocol Field" section under *Common Output Fields Description*.	**brief**

Table 16: show interfaces (Aggregated Ethernet) Output Fields *(continued)*

Field Name	Field Description	Level of Output
Protocol	Protocol family configured on the logical interface. Possible values are described in the "Protocol Field" section under *Common Output Fields Description*.	**detail extensive** none
MTU	Maximum transmission unit size on the logical interface.	**detail extensive** none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Routing table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive**
Flags	Information about protocol family flags. Possible values are described in the "Family Flags Field" section under *Common Output Fields Description*.	**detail extensive** none
Mac-Validate Failures	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	**detail extensive** none
Addresses, Flags	Information about address flags. Possible values are described in the "Addresses Flags" section under *Common Output Fields Description*.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces (Aggregated Ethernet)

```
user@host> show interfaces ae0
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 153, SNMP ifIndex: 59
  Link-level type: Ethernet, MTU: 1514, Speed: 300mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1
  Device flags   : Present Running
  Interface flags: SNMP-Traps 16384
  Current address: 00:05:85:8b:bf:f0, Hardware address: 00:05:85:8b:bf:f0
  Last flapped   : Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface ae0.0 (Index 72) (SNMP ifIndex 60)
    Flags: SNMP-Traps 16384 Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :             0          0             0            0
```

```
      Output:             0          0             0            0
  Protocol inet, MTU: 1500
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 10.100.1/24, Local: 10.100.1.2, Broadcast: 10.100.1.255
```

show interfaces brief (Aggregated Ethernet)

```
user@host> show interfaces ae0 brief
Physical interface: ae0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1514, Speed: 300mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps 16384

  Logical interface ae0.0
    Flags: SNMP-Traps 16384 Encapsulation: ENET2
    inet  10.100.1.2/24
```

show interfaces detail (Aggregated Ethernet)

```
user@host> show interfaces ae0 detail
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 153, SNMP ifIndex: 59, Generation: 36
  Link-level type: Ethernet, MTU: 1514, Speed: 300mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1
  Device flags   : Present Running
  Interface flags: SNMP-Traps 16384
  Current address: 00:05:85:8b:bf:f0, Hardware address: 00:05:85:8b:bf:f0
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                    0 bps
   Output bytes  :                   0                    0 bps
   Input  packets:                   0                    0 pps
   Output packets:                   0                    0 pps
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                7375                 7375                    0

    1 expedited-fo                  0                    0                    0

    2 assured-forw                  0                    0                    0

    3 network-cont               2268                 2268                    0


  Logical interface ae0.0 (Index 72) (SNMP ifIndex 60) (Generation 18)
    Flags: SNMP-Traps 16384 Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :             0          0             0            0
        Output:             0          0             0            0
    Link:
      fe-0/1/0.0
        Input :             0          0             0            0
        Output:             0          0             0            0
      fe-0/1/2.0
        Input :             0          0             0            0
        Output:             0          0             0            0
      fe-0/1/3.0
```

```
        Input :             0          0             0            0
        Output:             0          0             0            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      fe-0/1/0.0                0          0            0            0
      fe-0/1/2.0                0          0            0            0
      fe-0/1/3.0                0          0            0            0
    Protocol inet, MTU: 1500, Generation: 37, Route table: 0
      Flags: Is-Primary, Mac-Validate-Strict
      Mac-Validate Failures: Packets: 0, Bytes: 0
        Destination: 10.100.1/24, Local: 10.100.1.2, Broadcast: 10.100.1.255,
        Generation: 49
```

show interfaces extensive (Aggregated Ethernet)

```
user@host> show interfaces ae0 extensive
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 153, SNMP ifIndex: 59, Generation: 36
  Link-level type: Ethernet, MTU: 1514, Speed: 300mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1
  Device flags   : Present Running
  Interface flags: SNMP-Traps 16384
  Current address: 00:05:85:8b:bf:f0, Hardware address: 00:05:85:8b:bf:f0
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   60                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    1                    0 pps
   Output packets:                    0                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                 7375                 7375                    0

    1 expedited-fo                   0                    0                    0

    2 assured-forw                   0                    0                    0

    3 network-cont                2268                 2268                    0


  Logical interface ae0.0 (Index 72) (SNMP ifIndex 60) (Generation 18)
    Flags: SNMP-Traps 16384 Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :             1          0            60            0
        Output:             0          0             0            0
    Link:
      fe-0/1/0.0
        Input :             0          0             0            0
        Output:             0          0             0            0
      fe-0/1/2.0
        Input :             0          0             0            0
        Output:             0          0             0            0
      fe-0/1/3.0
        Input :             1          0            60            0
```

```
      Output:             0          0             0            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      fe-0/1/0.0                0           0            0            0
      fe-0/1/2.0                0           0            0            0
      fe-0/1/3.0                0           0            0            0
    Protocol inet, MTU: 1500, Generation: 37, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.100.1/24, Local: 10.100.1.2, Broadcast: 10.100.1.255,
        Generation: 49
```

show interfaces extensive (Aggregated Ethernet with VLAN Stacking)

```
user@host> show interfaces ae0 extensive
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 155, SNMP ifIndex: 48, Generation: 186
  Link-level type: 52, MTU: 1518, Speed: 2000mbps, Loopback: Disabled, Source
filtering: Disabled,
  Flow control: Disabled, Minimum links needed: 1, Minimum bandwidth needed: 0
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:12:1e:19:3f:f0, Hardware address: 00:12:1e:19:3f:f0
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              2406875                40152 bps
   Output bytes  :              1124470                22056 bps
   Input  packets:                 5307                    5 pps
   Output packets:                13295                   21 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
 0
  Ingress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                    0               859777                    0

    1 expedited-fo                   0                    0                    0

    2 assured-forw                   0                    0                    0

    3 network-cont                   0                    0                    0

  Egress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                    0              1897615                    0

    1 expedited-fo                   0                    0                    0

    2 assured-forw                   0                    0                    0

    3 network-cont                   0               662505                    0
```

```
  Logical interface ae0.451 (Index 69) (SNMP ifIndex 167) (Generation 601)
    Flags: SNMP-Traps VLAN-Tag [ 0x8100.451 ]  Encapsulation: VLAN-VPLS
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :           289          0         25685          376
        Output:          1698          4        130375         3096
    Link:
      ge-1/2/0.451
        Input :           289          0         25685          376
        Output:             0          0             0            0
      ge-1/2/1.451
        Input :             0          0             0            0
        Output:          1698          4        130375         3096
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      ge-1/2/0.451              0           0            0            0
      ge-1/2/1.451              0           0            0            0
    Protocol vpls, MTU: 1518, Generation: 849, Route table: 3
      Flags: Is-Primary

  Logical interface ae0.452 (Index 70) (SNMP ifIndex 170) (Generation 602)
    Flags: SNMP-Traps VLAN-Tag [ 0x8100.452 ]  Encapsulation: VLAN-VPLS
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :           293          1         26003         1072
        Output:          1694          3        130057         2400
    Link:
      ge-1/2/0.452
        Input :           293          1         26003         1072
        Output:          1694          3        130057         2400
      ge-1/2/1.452
        Input :             0          0             0            0
        Output:             0          0             0            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      ge-1/2/0.452              0           0            0            0
      ge-1/2/1.452              0           0            0            0
    Protocol vpls, MTU: 1518, Generation: 850, Route table: 3
      Flags: None
...
```

show interfaces (Fast Ethernet)

Syntax

show interfaces *interface-type*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display status information about the specified Fast Ethernet interface.

Options

interface-type—On M Series and T Series routers, the interface type is **fe-*fpc/pic/port***. On the J Series routers, the interface type is **fe-*pim*/0/*port***.

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

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

List of Sample Output

show interfaces (Fast Ethernet) on page 247
show interfaces brief (Fast Ethernet) on page 248
show interfaces detail (Fast Ethernet) on page 248
show interfaces extensive (Fast Ethernet) on page 248

Output Fields

Table 17 on page 234 lists the output fields for the **show interfaces Fast Ethernet** command. Output fields are listed in the approximate order in which they appear.

Table 17: show interfaces Fast Ethernet Output Fields

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

Table 17: show interfaces Fast Ethernet Output Fields *(continued)*

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

Table 17: show interfaces Fast Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Frequency	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels
CoS queues	Number of CoS queues configured.	**detail extensive** none
Schedulers	(GigabitEthernet intelligent queuing 2 (IQ2) interfaces only) Number of CoS schedulers configured.	**extensive**
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive** none
Hardware address	Hardware MAC address.	**detail extensive** none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour:minute:second:timezone*** **(*****hour:minute:second*** **ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).**	**detail extensive** none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface. Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes. For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 31 under the *show interfaces (10-Gigabit Ethernet)* command.	**detail extensive**

Table 17: 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 17: show interfaces Fast Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Collisions**—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • **Aged packets**—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • **FIFO errors**—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**
Egress queues	Total number of egress queues supported on the specified interface.	**detail extensive**
Queue counters (Egress)	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**
Ingress queues	Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.	**extensive**
Queue counters (Ingress)	CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**extensive**

Table 17: show interfaces Fast Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the routing device configuration, an alarm can ring the red or yellow alarm bell on the routing device, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value **None** or **Link**. • **None**—There are no active defects or alarms. • **Link**—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	**detail extensive** none
OTN FEC statistics	The forward error correction (FEC) counters provide the following statistics:. • **Corrected Errors**—The count of corrected errors in the last second. • **Corrected Error Ratio**—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits.	
PCS statistics	(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device. • **Bit errors**—High bit error rate. Indicates the number of bit errors when the PCS receiver is operating in normal mode. • **Errored blocks**—Loss of block lock. The number of errored blocks when PCS receiver is operating in normal mode.	**detail extensive**

Table 17: show interfaces Fast Ethernet Output Fields *(continued)*

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

Table 17: 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 17: show interfaces Fast Ethernet Output Fields *(continued)*

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

Table 17: show interfaces Fast Ethernet Output Fields *(continued)*

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

Table 17: show interfaces Fast Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output

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

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

Table 17: show interfaces Fast Ethernet Output Fields *(continued)*

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

Table 17: show interfaces Fast Ethernet Output Fields *(continued)*

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

Table 17: show interfaces Fast Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Preferred source address	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	**detail extensive** none
Input Filters	Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	**detail extensive**
Output Filters	Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	**detail extensive**
Mac-Validate Failures	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	**detail extensive** none
Addresses, Flags	Information about the address flags. Possible values are described in the "Addresses Flags" section under *Common Output Fields Description*.	**detail extensive** none
protocol-family	Protocol family configured on the logical interface. If the protocol is **inet**, the IP address of the interface is also displayed.	**brief**
Flags	Information about address flag (possible values are described in the "Addresses Flags" section under *Common Output Fields Description*.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interlace.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces (Fast Ethernet)

```
user@host> show interfaces fe-0/0/0
Physical interface: fe-0/0/0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 22
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:05:85:02:38:00, Hardware address: 00:05:85:02:38:00
  Last flapped   : 2006-01-20 14:50:58 PST (2w4d 00:44 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198)
    Flags: SNMP-Traps Encapsulation: ENET2
```

```
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255
```

show interfaces brief (Fast Ethernet)

```
user@host> show interfaces fe-0/0/0 brief
Physical interface: fe-0/0/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Logical interface fe-0/0/0.0
    Flags: SNMP-Traps Encapsulation: ENET2
    inet  10.10.10.1/24
```

show interfaces detail (Fast Ethernet)

```
user@host> show interfaces fe-0/0/0 detail
Physical interface: fe-0/0/0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 22, Generation: 5391
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:05:85:02:38:00, Hardware address: 00:05:85:02:38:00
  Last flapped   : 2006-01-20 14:50:58 PST (2w4d 00:45 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                  42                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                   1                   0 pps
  Active alarms  : None
  Active defects : None
  Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198) (Generation 67)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Generation: 105, Route table: 0
      Flags: Is-Primary, Mac-Validate-Strict
      Mac-Validate Failures: Packets: 0, Bytes: 0
      Addresses, Flags: Is-Preferred Is-Primary
      Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255,
        Generation: 136
```

show interfaces extensive (Fast Ethernet)

```
user@host> show interfaces fe-0/0/0 extensive
Physical interface: fe-0/0/0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 22, Generation: 5391
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed:
  100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:05:85:02:38:00, Hardware address: 00:05:85:02:38:00
  Last flapped   : 2006-01-20 14:50:58 PST (2w4d 00:46 ago)
  Statistics last cleared: Never
```

```
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                  42                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                   1                   0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 3, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Active alarms  : None
  Active defects : None
  MAC statistics:                      Receive         Transmit
    Total octets                             0               64
    Total packets                            0                1
    Unicast packets                          0                0
    Broadcast packets                        0                1
    Multicast packets                        0                0
    CRC/Align errors                         0                0
    FIFO errors                              0                0
    MAC control frames                       0                0
    MAC pause frames                         0                0
    Oversized frames                         0
    Jabber frames                            0
    Fragment frames                          0
    VLAN tagged frames                       0
    Code violations                          0
Filter statistics:
    Input packet count                       0
    Input packet rejects                     0
    Input DA rejects                         0
    Input SA rejects                         0
    Output packet count                                       1
    Output packet pad count                                   0
    Output packet error count                                 0
    CAM destination filters: 1, CAM source filters: 0
  Autonegotiation information:
    Negotiation status: Complete
    Link partner:
        Link partner: Full-duplex, Flow control: None, Remote fault: Ok
    Local resolution:
Packet Forwarding Engine configuration:
    Destination slot: 0
CoS information:
                 Bandwidth           Buffer Priority   Limit
                             %          bps   %         usec
    0 best-effort            95    950000000  95            0      low    none
    3 network-control         5     50000000   5            0      low    none
  Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198) (Generation 67)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Generation: 105, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255,
        Generation: 136
```

show interfaces (Gigabit Ethernet)

Syntax

show interfaces ge-*fpc/pic/port*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced before Junos OS Release 7.4.

Description

(M Series, T Series, and MX Series routers and EX Series switches only) Display status information about the specified Gigabit Ethernet interface.

Options

ge-*fpc/pic/port*—Display standard information about the specified Gigabit Ethernet interface.

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

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Additional Information

In a logical system, this command displays information only about the logical interfaces and not about the physical interfaces.

Required Privilege Level

view

Related Documentation

- *Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*

List of Sample Output

Output Fields

Table 18 on page 251 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 19 on page 264.

Table 18: show interfaces Gigabit Ethernet Output Fields

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

Table 18: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Link flags	Information about the link. Possible values are described in the "Links Flags" section under *Common Output Fields Description*.	All levels
Wavelength	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels
Frequency	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels
CoS queues	Number of CoS queues configured.	**detail extensive** none
Schedulers	(Gigabit Ethernet intelligent queuing 2 [IQ2] interfaces only) Number of CoS schedulers configured.	**extensive**
Hold-times	Current interface hold-time up and hold-time down, in milliseconds (ms).	**detail extensive**
Current address	Configured MAC address.	**detail extensive** none
Hardware address	Hardware MAC address.	**detail extensive** none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour:minute:second:timezone*** **(*****hour:minute:second*** **ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).**	**detail extensive** none
Input Rate	Input rate in bits per second (bps) and packets per second (pps). 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 accounting overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for egress traffic.	**detail extensive**
Ingress accounting overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for ingress traffic.	**detail extensive**

Table 18: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Output bytes**—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface. Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes. For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 31 under the *show interfaces (10-Gigabit Ethernet)* command.	**detail extensive**
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • **L3 incompletes**—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the **ignore-l3-incompletes** statement. • **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **FIFO errors**—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 18: 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 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. 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**

Table 18: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**—The count of corrected errors in the last second. • **Corrected Error Ratio**—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits.	**detail extensive**
PCS statistics	(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device. • **Bit errors**—High bit error rate. Indicates the number of bit errors when the PCS receiver is operating in normal mode. • **Errored blocks**—Loss of block lock. The number of errored blocks when the PCS receiver is operating in normal mode.	**detail extensive**

Table 18: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output

MAC statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC subsystem, including the following:
	• **Total octets** and **total packets**—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see Table 31 under the *show interfaces (10-Gigabit Ethernet)* command.
	• **Unicast packets, Broadcast packets,** and **Multicast packets**—Number of unicast, broadcast, and multicast packets.
	• **CRC/Align errors**—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).
	• **FIFO error**—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning.
	• **MAC control frames**—Number of MAC control frames.
	• **MAC pause frames**—Number of MAC control frames with **pause** operational code.
	• **Oversized frames**—There are two possible conditions regarding the number of oversized frames:
	  ◦ Packet length exceeds 1518 octets, or
	  ◦ Packet length exceeds MRU
	• **Jabber frames**—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.
	• **Fragment frames**—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.
	• **VLAN tagged frames**—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not.
	NOTE: The 20-port Gigabit Ethernet MIC (MIC-3D-20GE-SFP) does not have hardware counters for VLAN frames. Therefore, the **VLAN tagged frames** field displays 0 when the **show interfaces** command is executed on a 20-port Gigabit Ethernet MIC. In other words, the number of VLAN tagged frames cannot be determined for the 20-port Gigabit Ethernet MIC.
	• **Code violations**—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”
		extensive

OTN Received Overhead Bytes	APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08	**extensive**
OTN Transmitted Overhead Bytes	APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08	**extensive**

Table 18: show interfaces Gigabit Ethernet Output Fields *(continued)*

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

Table 18: show interfaces Gigabit Ethernet Output Fields *(continued)*

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

Table 18: 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 18: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output

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

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

Table 18: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface. • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP interface index number for the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under *Common Output Fields Description*.	All levels

Table 18: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
VLAN-Tag	Rewrite profile applied to incoming or outgoing frames on the outer (**Out**) VLAN tag or for both the outer and inner (**In**) VLAN tags. • **push**—An outer VLAN tag is pushed in front of the existing VLAN tag. • **pop**—The outer VLAN tag of the incoming frame is removed. • **swap**—The outer VLAN tag of the incoming frame is overwritten with the user-specified VLAN tag information. • **push**—An outer VLAN tag is pushed in front of the existing VLAN tag. • **push-push**—Two VLAN tags are pushed in from the incoming frame. • **swap-push**—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • **swap-swap**—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user-specified VLAN tag value. • **pop-swap**—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • **pop-pop**—Both the outer and inner VLAN tags of the incoming frame are removed.	**brief detail extensive** none
Demux	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: • Source Family Inet • Destination Family Inet	**detail extensive** none
Encapsulation	Encapsulation on the logical interface.	All levels
ACI VLAN: Dynamic Profile	Name of the dynamic profile that defines the agent circuit identifier (ACI) interface set. If configured, the ACI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ACI information.	**brief detail extensive** none
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under *Common Output Fields Description*.	**detail extensive** none
MTU	Maximum transmission unit size on the logical interface.	**detail extensive** none
Dynamic Profile	(MX Series routers with Trio MPCs only) Name of the dynamic profile that was used to create this interface configured with a Point-to-Point Protocol over Ethernet (PPPoE) family.	**detail extensive** none
Service Name Table	(MX Series routers with Trio MPCs only) Name of the service name table for the interface configured with a PPPoE family.	**detail extensive** none
Max Sessions	(MX Series routers with Trio MPCs only) Maximum number of PPPoE logical interfaces that can be activated on the underlying interface.	**detail extensive** none

Table 18: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Duplicate Protection	(MX Series routers with Trio MPCs only) State of PPPoE duplicate protection: **On** or **Off**. When duplicate protection is configured for the underlying interface, a dynamic PPPoE logical interface cannot be activated when an existing active logical interface is present for the same PPPoE client.	**detail extensive** none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	**detail extensive** none
Traffic statistics	Number and rate of bytes and packets received and transmitted on the specified interface set. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Input packets, Output packets**—Number of packets received and transmitted on the interface set.	**detail extensive**
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	**extensive**
Local statistics	Number and rate of bytes and packets destined to the router.	**extensive**
Transit statistics	Number and rate of bytes and packets transiting the switch. NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the **Output bytes** and **Output packets** interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.	**extensive**
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Route table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive** none
Flags	Information about protocol family flags. Possible values are described in the "Family Flags" section under *Common Output Fields Description*.	**detail extensive**
Donor interface	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	**detail extensive** none
Preferred source address	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	**detail extensive** none
Input Filters	Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parentheses next to all interfaces.	**detail extensive**

Table 18: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Output Filters	Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parentheses next to all interfaces.	**detail extensive**
Mac-Validate Failures	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	**detail extensive** none
Addresses, Flags	Information about the address flags. Possible values are described in the "Addresses Flags" section under *Common Output Fields Description*.	**detail extensive** none
protocol-family	Protocol family configured on the logical interface. If the protocol is **inet**, the IP address of the interface is also displayed.	**brief**
Flags	Information about the address flag. Possible values are described in the "Addresses Flags" section under *Common Output Fields Description*.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Table 19: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type

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

Table 19: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type *(continued)*

Interface Type	Sample Command	Byte and Octet Counts Include	Comments
Outbound logical interface	**show interfaces ge-0/0/0.50 extensive**	Traffic statistics: Input bytes: 478 bytes per packet, representing the Layer 3 packet	

Sample Output

show interfaces (Gigabit Ethernet)

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

  Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69)
    Flags: SNMP-Traps 0x4000
    VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push
    0x8100.512 0x8100.513)
    Encapsulation: VLAN-CCC
    Egress accounting overhead: 100
    Ingress accounting overhead: 90
    Input packets : 0
    Output packets: 0
    Protocol ccc, MTU: 1522
      Flags: Is-Primary
```

show interfaces (Gigabit Ethernet on MX Series Routers)

```
user@host> show interfaces ge-2/2/2
Physical interface: ge-2/2/2, Enabled, Physical link is Up
  Interface index: 156, SNMP ifIndex: 188
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, MAC-REWRITE Error: None,
 Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Schedulers     : 0
  Current address: 00:1f:12:b7:d7:c0, Hardware address: 00:1f:12:b7:d6:76
  Last flapped   : 2008-09-05 16:44:30 PDT (3d 01:04 ago)
  Input rate     : 0 bps (0 pps)
```

```
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  Logical interface ge-2/2/2.0 (Index 82) (SNMP ifIndex 219)
    Flags: SNMP-Traps 0x20000000 Encapsulation: Ethernet-Bridge
    Egress accounting overhead: 100
    Ingress accounting overhead: 90
    Input packets : 0
    Output packets: 0
     Protocol aenet, AE bundle: ae0.0    Link Index: 4
```

show interfaces extensive (Gigabit Ethernet on MX Series Routers showing interface transmit 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
```

show interfaces brief (Gigabit Ethernet)

```
user@host> show interfaces ge-3/0/2 brief
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface ge-3/0/2.0
    Flags: SNMP-Traps 0x4000
    VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push
    0x8100.512 0x8100.513)
    Encapsulation: VLAN-CCC
    ccc

  Logical interface ge-3/0/2.32767
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
```

show interfaces detail (Gigabit Ethernet)

```
user@host> show interfaces ge-3/0/2 detail
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Interface index: 167, SNMP ifIndex: 35, Generation: 177
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
```

```
  Current address: 00:05:85:4a:e9:7c, Hardware address: 00:05:85:4a:e9:7c
  Last flapped   : 2006-08-09 17:17:00 PDT (01:31:33 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                   0                   0 pps
  Ingress traffic statistics at Packet Forwarding Engine:
   Input  bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Drop   bytes  :                   0                   0 bps
   Drop   packets:                   0                   0 pps
  Ingress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                   0                   0                    0

    1 expedited-fo                  0                   0                    0

    2 assured-forw                  0                   0                    0

    3 network-cont                  0                   0                    0

  Egress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                   0                   0                    0

    1 expedited-fo                  0                   0                    0

    2 assured-forw                  0                   0                    0

    3 network-cont                  0                   0                    0

  Active alarms  : None
  Active defects : None

  Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69) (Generation 140)
    Flags: SNMP-Traps 0x4000
    VLAN-Tag [0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530)
Out(swap-push 0x8100.512 0x8100.513)
    Encapsulation: VLAN-CCC
    Egress accounting overhead: 100
    Ingress accounting overhead: 90
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :                   0                   0 bps
     Output bytes  :                   0                   0 bps
     Input  packets:                   0                   0 pps
     Output packets:                   0                   0 pps
    Protocol ccc, MTU: 1522, Generation: 149, Route table: 0
```

```
    Flags: Is-Primary

  Logical interface ge-3/0/2.32767 (Index 71) (SNMP ifIndex 70)
  (Generation 139)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :                   0                    0 bps
     Output bytes  :                   0                    0 bps
     Input  packets:                   0                    0 pps
     Output packets:                   0                    0 pps
```

show interfaces extensive (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: 00:14:f6:30:5e:74, Hardware address: 00:14:f6:30:5e:74
  Last flapped   : 2007-11-07 21:31:41 PST (02:03:33 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :          38910844056                 7952 bps
   Output bytes  :              7174605                 8464 bps
   Input  packets:            418398473                   11 pps
   Output packets:                78903                   12 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Ingress traffic statistics at Packet Forwarding Engine:
   Input  bytes  :          38910799145                 7952 bps
   Input  packets:            418397956                   11 pps
   Drop   bytes  :                   0                    0 bps
   Drop   packets:                   0                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
```

```
Ingress queues: 4 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets

  0 best-effort            418390823            418390823                    0

  1 expedited-fo                   0                    0                    0

  2 assured-forw                   0                    0                    0

  3 network-cont                7133                 7133                    0

Egress queues: 4 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets

  0 best-effort                 1031                 1031                    0

  1 expedited-fo                   0                    0                    0

  2 assured-forw                   0                    0                    0

  3 network-cont               77872                77872                    0

Active alarms  : None
Active defects : None
MAC statistics:                      Receive         Transmit
  Total octets                   38910844056          7174605
  Total packets                    418398473            78903
  Unicast packets                408021893366             1026
  Broadcast packets                       10               12
  Multicast packets                418398217            77865
  CRC/Align errors                         0                0
  FIFO errors                              0                0
  MAC control frames                       0                0
  MAC pause frames                         0                0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0  OTN Received Overhead Bytes:
  APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58
  Payload Type: 0x08
OTN Transmitted Overhead Bytes:
  APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
  Payload Type: 0x08
Filter statistics:
  Input packet count               418398473
  Input packet rejects                   479
  Input DA rejects                       479
  Input SA rejects                         0
  Output packet count                                   78903
  Output packet pad count                                   0
  Output packet error count                                 0
  CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Complete
  Link partner:
      Link mode: Full-duplex, Flow control: Symmetric/Asymmetric,
      Remote fault: OK
  Local resolution:
      Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
```

```
    Destination slot: 7
  CoS information:
    Direction : Output
    CoS transmit queue          Bandwidth          Buffer     Priority   Limit
                              %            bps     %           usec
    0 best-effort            95      950000000    95              0
low    none
    3 network-control         5       50000000     5              0
low    none
    Direction : Input
    CoS transmit queue          Bandwidth          Buffer     Priority   Limit
                              %            bps     %           usec
    0 best-effort            95      950000000    95              0
low    none
    3 network-control         5       50000000     5              0
low    none

  Logical interface ge-7/1/3.0 (Index 70) (SNMP ifIndex 85) (Generation 150)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :               812400
     Output bytes  :              1349206
     Input  packets:                 9429
     Output packets:                 9449
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Local statistics:
     Input  bytes  :               812400
     Output bytes  :              1349206
     Input  packets:                 9429
     Output packets:                 9449
    Transit statistics:
     Input  bytes  :                   0                 7440 bps
     Output bytes  :                   0                 7888 bps
     Input  packets:                   0                   10 pps
     Output packets:                   0                   11 pps
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Protocol inet, MTU: 1500, Generation: 169, Route table: 0
      Flags: Is-Primary, Mac-Validate-Strict
      Mac-Validate Failures: Packets: 0, Bytes: 0
      Addresses, Flags: Is-Preferred Is-Primary
      Input Filters: F1-ge-3/0/1.0-in, F3-ge-3/0/1.0-in
      Output Filters: F2-ge-3/0/1.0-out (53)
      Destination: 10.74.2/24, Local: 10.74.2.2, Broadcast: 10.74.2.255,
        Generation: 196
    Protocol multiservice, MTU: Unlimited, Generation: 170, Route table: 0
      Flags: Is-Primary
      Policer: Input: __default_arp_policer__
```

NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics displayed in the **show interfaces** command output might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the interface counters. For detailed

information, see the description of the logical interface **Transit statistics** fields in Table 18 on page 251.

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: 00:14:f6:11:26:f8, Hardware address: 00:14:f6:11:26:f8
  Last flapped   : 2006-10-27 04:42:23 PDT (08:01:52 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 624 bps (1 pps)
  Active alarms  : None
  Active defects : None

  Logical interface ge-3/2/0.0 (Index 67) (SNMP ifIndex 85)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 0
    Output packets: 6
    Protocol inet, MTU: 1500
      Flags: Unnumbered
      Donor interface: lo0.0 (Index 64)
      Preferred source address: 22.22.22.22
```

show interfaces (ACI Interface Set Configured)

```
user@host> show interfaces ge-1/0/0.4001
  Logical interface ge-1/0/0.4001 (Index 340) (SNMP ifIndex 548)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.4001 ]  Encapsulation: PPP-over-

    Ethernet
    ACI VLAN:
      Dynamic Profile: aci-vlan-set-profile
    PPPoE:
      Dynamic Profile: aci-vlan-pppoe-profile,
      Service Name Table: None,
      Max Sessions: 32000, Max Sessions VSA Ignore: Off,
      Duplicate Protection: On, Short Cycle Protection: Off,
      AC Name: nbc
    Input packets : 9
    Output packets: 8
    Protocol multiservice, MTU: Unlimited
```

show interfaces demux0 (Demux Interfaces)

Syntax

show interfaces demux0.*logical-interface-number*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced in Junos OS Release 9.0.

Description

(MX Series and M Series routers only) Display status information about the specified demux interface.

Options

none—Display standard information about the specified demux interface.

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

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Related Documentation

- *Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*

List of Sample Output

show interfaces (Demux) on page 278
show interfaces (PPPoE over Aggregated Ethernet) on page 279
show interfaces extensive (Targeted Distribution for Aggregated Ethernet Links) on page 279
show interfaces demux0 (ACI Interface Set Configured) on page 280

Output Fields

Table 20 on page 272 lists the output fields for the **show interfaces** (demux interfaces) command. Output fields are listed in the approximate order in which they appear.

Table 20: Demux show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	**brief detail extensive** none
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**brief detail extensive** none

Table 20: Demux show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Enabled	State of the interface. Possible values are described in the "Enabled Field" section under *Common Output Fields Description*.	**brief detail extensive** none
Physical link	Status of the physical link (**Up** or **Down**).	**detail extensive** none
Admin	Administrative state of the interface (**Up** or **Down**).	**terse**
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**detail extensive** none
Link	Status of the physical link (**Up** or **Down**).	**terse**
Targeting summary	Status of aggregated Ethernet links that are configured with targeted distribution (**primary** or **backup**)	**extensive**
Bandwidth	Bandwidth allocated to the aggregated Ethernet links that are configured with targeted distribution.	**extensive**
Proto	Protocol family configured on the interface.	**terse**
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Type	Type of interface. **Software-Pseudo** indicates a standard software interface with no associated hardware device.	**brief detail extensive** none
Link-level type	Encapsulation being used on the physical interface.	**brief detail extensive**
MTU	Maximum transmission unit size on the physical interface.	**brief detail extensive**
Clocking	Reference clock source: **Internal** (1) or **External** (2).	**brief detail extensive**
Speed	Speed at which the interface is running.	**brief detail extensive**
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under *Common Output Fields Description*.	**brief detail extensive** none
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under *Common Output Fields Description*.	**brief detail extensive** none
Link type	Data transmission type.	**detail extensive** none
Link flags	Information about the link. Possible values are described in the "Link Flags" section under *Common Output Fields Description*.	**detail extensive** none
Physical info	Information about the physical interface.	**detail extensive**
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**

Table 20: Demux show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Current address	Configured MAC address.	**detail extensive**
Hardware address	Hardware MAC address.	**detail extensive**
Alternate link address	Backup address of the link.	**detail extensive**
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour:minute:second:timezone*** **(*****hour:minute:second*** **ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).**	**detail extensive** none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**

Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface.
	• **Input bytes**—Number of bytes received on the interface.
	• **Output bytes**—Number of bytes transmitted on the interface.
	• **Input packets**—Number of packets received on the interface.
	• **Output packets**—Number of packets transmitted on the interface.
	• **IPv6 transit statistics**—Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled.
	NOTE: These fields include dropped traffic and exception traffic, as those fields are not separately defined.
	  • **Input bytes**—Number of bytes received on the interface
	  • **Output bytes**—Number of bytes transmitted on the interface.
	  • **Input packets**—Number of packets received on the interface.
	  • **Output packets**—Number of packets transmitted on the interface.
		detail extensive

Input errors	Input errors on the interface whose definitions are as follows: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Giants**—Number of frames received that are larger than the giant packet threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • **Resource errors**—Sum of transmit drops.	**extensive**
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	none

Table 20: Demux show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**
Output Rate	Output rate in bps and pps.	none
Logical Interface		
Logical interface	Name of the logical interface.	**brief detail extensive** none
Index	Index number of the logical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP interface index number for the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail**
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under *Common Output Fields Description*.	**brief detail extensive** none
Encapsulation	Encapsulation on the logical interface.	**brief extensive** none
ACI VLAN: Dynamic Profile	Name of the dynamic profile that defines the agent circuit identifier (ACI) interface set. If configured, the ACI interface set enables the underlying demux interface to create dynamic VLAN subscriber interfaces based on ACI information.	**brief detail extensive** none
Demux	Specific IP demultiplexing (demux) values: • **Underlying interface**—The underlying interface that the demux interface uses. • **Index**—Index number of the logical interface. • **Family**—Protocol family configured on the logical interface. • **Source prefixes, total**—Total number of source prefixes for the underlying interface. • **Destination prefixes, total**—Total number of destination prefixes for the underlying interface. • **Prefix**—**inet** family prefix.	**detail extensive** none

Table 20: Demux show interfaces Output Fields *(continued)*

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

Table 20: Demux show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Output packets	Number of packets transmitted on the interface.	none
Protocol	Protocol family. Possible values are described in the "Protocol Field" section under *Common Output Fields Description*.	**detail extensive** none
MTU	Maximum transmission unit size on the logical interface.	**detail extensive** none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route table	Route table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive**
Flags	Information about protocol family flags. Possible values are described in the "Family Flags" section under *Common Output Fields Description*.	**detail extensive** none
Mac-Validate Failures	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	**detail extensive** none
Addresses, Flags	Information about the address flags. Possible values are described in the "Addresses Flags" section under *Common Output Fields Description*.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive statistics** none
Local	IP address of the logical interface.	**detail extensive terse** none
Remote	IP address of the remote interface.	**terse**
Broadcast	Broadcast address of the logical interlace.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Link	Name of the physical interfaces for member links in an aggregated Ethernet bundle for a PPPoE over aggregated Ethernet configuration. PPPoE traffic goes out on these interfaces.	**detail extensive** none
Dynamic-profile	Name of the PPPoE dynamic profile assigned to the underlying interface.	**detail extensive** none
Service Name Table	Name of the PPPoE service name table assigned to the PPPoE underlying interface.	**detail extensive** none
Max Sessions	Maximum number of dynamic PPPoE logical interfaces that the router can activate on the underlying interface.	**detail extensive** none

Table 20: Demux show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Duplicate Protection	State of duplicate protection: **On** or **Off**. Duplicate protection prevents the activation of another dynamic PPPoE logical interface on the same underlying interface when a dynamic PPPoE logical interface for a client with the same MAC address is already active on that interface.	**detail extensive** none
AC Name	Name of the access concentrator.	**detail extensive** none

Sample Output

show interfaces (Demux)

```
user@host> show interfaces demux0
Physical interface: demux0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 79, Generation: 129
  Type: Software-Pseudo, Link-level type: Unspecified, MTU: 9192, Clocking: 1,
  Speed: Unspecified
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                   0                   0 pps
   IPv6 transit statistics:
    Input  bytes  :                  0
    Output bytes  :                  0
    Input  packets:                  0
    Output packets:                  0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0

Logical interface demux0.0 (Index 87) (SNMP ifIndex 84) (Generation 312)
    Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
    Demux:
     Underlying interface: ge-2/0/1.0 (Index 74)
   Family Inet Source prefixes, total 1
   Prefix: 1.1.1/24
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                1554
     Input  packets:                   0
     Output packets:                  37
     IPv6 transit statistics:
```

```
      Input  bytes  :                    0
      Output bytes  :                    0
      Input  packets:                    0
      Output packets:                    0
    Local statistics:
     Input  bytes  :                     0
     Output bytes  :                  1554
     Input  packets:                     0
     Output packets:                    37
    Transit statistics:
     Input  bytes  :                     0                    0 bps
     Output bytes  :                     0                    0 bps
     Input  packets:                     0                    0 pps
     Output packets:                     0                    0 pps
     IPv6 transit statistics:
      Input  bytes  :                    0
      Output bytes  :                    0
      Input  packets:                    0
      Output packets:                    0
    Protocol inet, MTU: 1500, Generation: 395, Route table: 0
      Flags: Is-Primary, Mac-Validate-Strict
      Mac-Validate Failures: Packets: 0, Bytes: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 11.1.1/24, Local: 11.1.1.1, Broadcast: 11.1.1.255,
        Generation: 434
```

show interfaces (PPPoE over Aggregated Ethernet)

```
user@host> show interfaces demux0.100
Logical interface demux0.100 (Index 76) (SNMP ifIndex 61160)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ]
    Encapsulation: ENET2
    Demux:
      Underlying interface: ae0 (Index 199)
    Link:
      ge-1/0/0
      ge-1/1/0
    Input packets : 0
    Output packets: 0
    Protocol pppoe
      Dynamic Profile: pppoe-profile,
      Service Name Table: service-table1,
      Max Sessions: 100, Duplicate Protection: On,
      AC Name: pppoe-server-1
```

show interfaces extensive (Targeted Distribution for Aggregated Ethernet Links)

```
user@host> show interfaces demux0.1073741824 extensive

  Logical interface demux0.1073741824 (Index 75) (SNMP ifIndex 558) (Generation
346)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ]  Encapsulation: ENET2
    Demux:
      Underlying interface: ae0 (Index 201)
    Link:
      ge-1/0/0
      ge-1/1/0
      ge-2/0/7
      ge-2/0/8
    Targeting summary:
      ge-1/1/0, primary, Physical link is Up
```

```
      ge-2/0/8, backup, Physical link is Up
    Bandwidth: 1000mbps
```

show interfaces demux0 (ACI Interface Set Configured)

```
user@host> show interfaces demux0.1073741827
  Logical interface demux0.1073741827 (Index 346) (SNMP ifIndex 527)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1802 0x8100.302 ]  Encapsulation:
 ENET2
    Demux: Source Family Inet
    ACI VLAN:
      Dynamic Profile: aci-vlan-set-profile
    Demux:
      Underlying interface: ge-1/0/0 (Index 138)
    Input packets : 18
    Output packets: 16
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re, Unnumbered
      Donor interface: lo0.0 (Index 322)
      Preferred source address: 100.20.200.202
      Addresses, Flags: Primary Is-Default Is-Primary
        Local: 10.4.12.119
    Protocol pppoe
      Dynamic Profile: aci-vlan-pppoe-profile,
      Service Name Table: None,
      Max Sessions: 32000, Max Sessions VSA Ignore: Off,
      Duplicate Protection: On, Short Cycle Protection: Off,
      AC Name: nbc
```

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

(J Series Services Routers, 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 287
show interfaces (PPPoE over Aggregated Ethernet) on page 287
show interfaces brief (PPPoE) on page 288
show interfaces detail (PPPoE) on page 288
show interfaces detail (PPPoE on J Series Services Routers) on page 289
show interfaces extensive (PPPoE on M120 and M320 Routers) on page 290

Output Fields

Table 21 on page 281 lists the output fields for the **show interfaces** (PPPoE) command. Output fields are listed in the approximate order in which they appear.

Table 21: 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*.	All levels
Interface index	Physical interface index number, which reflects its initialization sequence.	**detail extensive** none

Table 21: 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*.	All levels
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under *Common Output Fields Description*.	All levels
Link 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*.	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 21: 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		

Table 21: 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*.	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 21: 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

Table 21: 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*.	**detail extensive** none

Table 21: 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*.	**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: 00:10:94:00:00:01,
      Underlying interface: demux0.100 (Index 88)
    Link:
      ge-1/0/0.32767
      ge-1/0/1.32767
    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: 45.63.24.1
```

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 detail (PPPoE on J Series Services Routers)

```
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
  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 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: 00:17:cb:48:c8:12,
      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
     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: 12.12.12.2, Local: 12.12.12.1, Broadcast: Unspecified,
Generation: 206
```

show interfaces filters

Syntax

show interfaces filters
<interface-name>

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 all firewall filters that are installed on each interface in a system.

Options

none—Display filter information about all interfaces.

interface-name—(Optional) Display filter information about a particular interface.

Additional Information

For information about how to configure firewall filters, see the *Routing Policy Feature Guide for Routing Devices*. For related operational mode commands, see the *Junos OS Operational Mode Commands*.

Required Privilege Level

view

List of Sample Output

show interfaces filters on page 293
show interfaces filters interface-name on page 293
show interfaces filters (PTX Series Packet Transport Routers) on page 293

Output Fields

Table 22 on page 292 lists the output fields for the **show interfaces filters** command. Output fields are listed in the approximate order in which they appear.

Table 22: show interfaces filters Output Fields

Field Name	Field Description
Interface	Name of the interface.
Admin	Interface state: **up** or **down**.
Link	Link state: **up** or **down**.
Proto	Protocol configured on the interface.
Input Filter	Names of any firewall filters to be evaluated when packets are received on the interface, including any filters attached through activation of dynamic service.
Output Filter	Names of any firewall filters to be evaluated when packets are transmitted on the interface, including any filters attached through activation of dynamic service.

Sample Output

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 filters interface-name

```
user@host> show interfaces filters so-2/1/0
Interface       Admin Link Proto Input Filter         Output Filter
so-2/1/0        up    down
so-2/1/0.0      up    down inet  goop                 outfilter
                           iso
                           inet6 v6in                 v6out
```

```
user@host > show interfaces filters ge-3/0/1
Interface       Admin Link Proto Input Filter         Output Filter
ge-3/0/1        up    up
ge-3/0/1.0      up    up   inet  F1-ge-3/0/1.0-in     F2-ge-3/0/1.0-out
                           inet  F3-ge-3/0/1.0-in
```

show interfaces filters (PTX Series Packet Transport Routers)

```
user@host > show interfaces filters em0
Interface       Admin Link Proto Input Filter         Output Filter
      em0             up    up
      em0.0           up    up   inet
```

show interfaces routing

Syntax

show interfaces routing
<brief | detail>
<*interface-name*>
<logical-system (all | *logical-system-name*)>

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display the state of the router's interfaces. Use this command for performing router diagnostics only, when you are determining whether the routing protocols and the Junos OS differ about the state of an interface.

Options

none—Display standard information about the state of all router interfaces on all logical systems.

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

interface-name—(Optional) Name of 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

For information about how to configure routing protocols, see the *Junos OS Routing Protocols Library for Routing Devices*. For information about related operational mode commands for routing instances and protocols, see the *Junos OS Operational Mode Commands*.

Required Privilege Level

view

List of Sample Output

show interfaces routing brief on page 295
show interfaces routing brief (TX Matrix Plus Router) on page 296
show interfaces routing detail on page 296
show interfaces routing detail (TX Matrix Plus Router) on page 297

Output Fields

Table 23 on page 294 lists the output fields for the **show interfaces routing** command. Output fields are listed in the approximate order in which they appear.

Table 23: show interfaces routing Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the physical interface.	none **brief**
State	State of the physical interface: **Up** or **Down**.	none **brief**
Addresses	Protocols and addresses configured on the interface.	none **brief**
Index	Interface index number, which reflects its initialization sequence.	**detail**

Table 23: show interfaces routing Output Fields *(continued)*

Field Name	Field Description	Level of Output
Refcount	Number of references to the interface in the routing software.	**detail**
State	State (**Up** or **Down**) and type of interface.	**detail**
Change	Reflects one or more of the following recent changes to the interface: • **Add**—The interface was just added. • **Address**—The interface's link-layer address has changed. • **Delete**—The interface is being deleted. • **Encapsulation**—The type of encapsulation on the interface has changed. • **Metric**—The interface's metric value has changed. • **MTU**—The interface's maximim transmission unit size has changed. • **UpDown**—The interface has made an up or down transition.	**detail**
Up/down transitions	Number of times the interface has gone from **Down** to **Up**.	**detail**
Link layer	Describes the link layer of the interface.	**detail**
Encapsulation	Encapsulation on the interface.	**detail**
Bandwidth	Speed at which the interface is running.	**detail**

***Protocol* address**	Information about the configuration of protocols on the interface:
	• **Address**—Address configured on the interface for the protocol type.
	• **State**—State (**Up** or **down**) and type of interface.
	• **Change**—Reflects one or more of the following recent changes to the interface:
	  ◦ **Add**—The interface was just added.
	  ◦ **Address**—The interface's address has changed.
	  ◦ **Broadcast**—The interface's broadcast address has changed.
	  ◦ **Delete**—The interface is being deleted.
	  ◦ **Netmask**—The interface's netmask has changed.
	  ◦ **UpDown**—The interface has made an up or down transition.
	• **Preference**—Preference value for the route for this address.
	• **Metric**—Metric value on the interface for the protocol type.
	• **MTU**—Maximim transmission unit value of the interface.
	• **Local address**—On a point-to-point link, the address of the local side of the link. Not used for multicast links.
	• **Destination**—For a point-to-point link, the address of the remote side of the link. For multicast links, the network address.
		detail

Sample Output

show interfaces routing brief

```
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 brief (TX Matrix Plus Router)

```
user@host> show interfaces routing brief
Interface        State Addresses
...
ge-23/0/4.0      Up    INET  2.9.1.1
                       ISO   enabled
                       MPLS  enabled
ge-23/0/3.0      Up    INET  2.8.1.1
                       ISO   enabled
                       MPLS  enabled
ge-23/0/2.0      Up    INET  2.7.1.1
                       ISO   enabled
                       MPLS  enabled
ge-23/0/1.0      Up    INET  2.6.1.1
                       ISO   enabled
                       MPLS  enabled
ge-23/0/0.0      Up    INET  2.5.1.1
                       ISO   enabled
                       MPLS  enabled
ge-31/0/7.599    Up    INET  2.14.10.93
ge-31/0/7.598    Up    INET  2.14.10.89
ge-31/0/7.597    Up    INET  2.14.10.85
ge-31/0/7.596    Up    INET  2.14.10.81
ge-31/0/7.595    Up    INET  2.14.10.77
ge-31/0/7.594    Up    INET  2.14.10.73
...
ixgbe1.0         Up    INET  10.34.0.4
                       INET  162.0.0.4
                       INET6 fe80::200:1ff:fe22:4
                       INET6 fec0::a:22:0:4
ixgbe0.0         Up    INET  10.34.0.4
                       INET  162.0.0.4
                       INET6 fe80::200:ff:fe22:4
                       INET6 fec0::a:22:0:4
em0.0            Up    INET  192.168.178.11
```

show interfaces routing detail

```
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 detail (TX Matrix Plus Router)

```
user@host> show interfaces routing detail
ge-23/0/4.0
  Index: 77, Refcount: 5, State: <Up Broadcast Multicast> Change: <>
  0 metric, 0 up/down transitions, reth state 0, full-duplex
  Link layer: Ethernet Encapsulation: Ethernet Bandwidth: 1000Mbps
  Link address #0 0.1d.b5.14.da.2d
  INET address 2.9.1.1
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <RT-Change>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 2.9.1.3
    Destination: 2.9.1.0/30
    System flags: <Is-Preferred Is-Primary>
  ISO address (null)
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1497 bytes
    System flags: <>
  MPLS address (null)
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1488 bytes
    System flags: <>
ge-23/0/3.0
  Index: 76, Refcount: 5, State: <Up Broadcast Multicast> Change: <>
  0 metric, 0 up/down transitions, reth state 0, full-duplex
  Link layer: Ethernet Encapsulation: Ethernet Bandwidth: 1000Mbps
  Link address #0 0.1d.b5.14.da.2c
  INET address 2.8.1.1
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <RT-Change>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 2.8.1.3
    Destination: 2.8.1.0/30
```

```
    System flags: <Is-Preferred Is-Primary>
  ISO address (null)
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1497 bytes
    System flags: <>
  MPLS address (null)
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1488 bytes
    System flags: <>
ge-23/0/2.0
  Index: 75, Refcount: 5, State: <Up Broadcast Multicast> Change: <>
  0 metric, 0 up/down transitions, reth state 0, full-duplex
  Link layer: Ethernet Encapsulation: Ethernet Bandwidth: 1000Mbps
  Link address #0 0.1d.b5.14.da.2b
  INET address 2.7.1.1
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <RT-Change>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 2.7.1.3
    Destination: 2.7.1.0/30
    System flags: <Is-Preferred Is-Primary>
  ISO address (null)
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1497 bytes
    System flags: <>
  MPLS address (null)
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1488 bytes
    System flags: <>
ge-23/0/1.0
  Index: 74, Refcount: 5, State: <Up Broadcast Multicast> Change: <>
  0 metric, 0 up/down transitions, reth state 0, full-duplex
  Link layer: Ethernet Encapsulation: Ethernet Bandwidth: 1000Mbps
  Link address #0 0.1d.b5.14.da.2a
  INET address 2.6.1.1
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <RT-Change>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 2.6.1.3
...
ixgbe1.0
  Index: 5, Refcount: 5, State: <Up Broadcast Multicast> Change: <>
  0 metric, 0 up/down transitions, reth state 0, full-duplex
  Link layer: Ethernet Encapsulation: Ethernet Bandwidth: 1000Mbps
  Link address #0 2.0.1.22.0.4
  INET address 10.34.0.4
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 10.255.255.255
    Destination: 10.0.0.0/8
    System flags: <Is-Preferred>
  INET address 162.0.0.4
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 191.255.255.255
    Destination: 128.0.0.0/2
    System flags: <Primary Is-Preferred Is-Primary>
  INET6 address fe80::200:1ff:fe22:4
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1500 bytes
    Destination: fe80::/64
    System flags: <Is-Preferred>
  INET6 address fec0::a:22:0:4
```



```
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1500 bytes
    Destination: fec0::/64
    System flags: <Is-Preferred Is-Primary>
ixgbe0.0
  Index: 4, Refcount: 5, State: <Up Broadcast Multicast> Change: <>
  0 metric, 0 up/down transitions, reth state 0, full-duplex
  Link layer: Ethernet Encapsulation: Ethernet Bandwidth: 1000Mbps
  Link address #0 2.0.0.22.0.4
  INET address 10.34.0.4
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 10.255.255.255
    Destination: 10.0.0.0/8
    System flags: <Is-Preferred>
  INET address 162.0.0.4
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 191.255.255.255
    Destination: 128.0.0.0/2
    System flags: <Primary Is-Default Is-Preferred Is-Primary>
  INET6 address fe80::200:ff:fe22:4
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1500 bytes
    Destination: fe80::/64
    System flags: <Is-Preferred>
  INET6 address fec0::a:22:0:4
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1500 bytes
    Destination: fec0::/64
    System flags: <Is-Default Is-Preferred Is-Primary>
em0.0
  Index: 3, Refcount: 2, State: <Up Broadcast Multicast> Change: <>
  0 metric, 0 up/down transitions, reth state 0, full-duplex
  Link layer: Ethernet Encapsulation: Ethernet Bandwidth: 100Mbps
  Link address #0 0.80.f9.26.0.c0
  INET address 192.168.178.11
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 192.168.178.127
    Destination: 192.168.178.0/25
    System flags: <Is-Preferred Is-Primary>
```

show ppp interface

Syntax show ppp interface *interface-name*
<extensive| terse>

Release Information Command introduced in Junos OS Release 7.5.

Description Display information about PPP interfaces.

Options ***interface-name***—Name of a logical interface.

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

Required Privilege Level view

List of Sample Output
show ppp interface on page 307
show ppp interface extensive on page 307
show ppp interface terse on page 308

Output Fields Table 24 on page 300 lists the output fields for the **show ppp interface** command. Output fields are listed in the approximate order in which they appear.

Table 24: show ppp interface Output Fields

Field Name	Field Description	Level of Output
Session	Name of the logical interface on which the session is running.	All levels
Type	Session type: PPP.	All levels
Phase	PPP process phase: **Authenticate**, **Pending**, **Establish**, **LCP**, **Network**, **Disabled**, and **Tunneled**.	All levels
Session flags	Special conditions present in the session: **Bundled**, **TCC**, **No-keepalives**, **Looped**, **Monitored**, and **NCP-only**.	All levels
***protocol* State**	Protocol state information. See specific protocol state fields for information.	None specified
AUTHENTICATION	Challenge-Handshake Authentication Protocol (CHAP) authentication state information or Password Authentication Protocol (PAP) state information. See the **Authentication** field description for further information.	None specified

Table 24: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
Keepalive settings	Keepalive settings for the PPP sessions on the L2TP network server (LNS). LNS based PPP sessions are supported only on service interfaces (si). • **Interval**—Time in seconds between successive keepalive requests. Keepalive aging timeout is calculated as a product of the **interval** and **Down-count** values. If the keepalive aging timeout is greater than 180 seconds, the keepalive packets are handled by the Routing Engine. If the aging timeout is less than or equal to 180 seconds, the packets are handled by the Packet Forwarding Engine. • **Up-count**—The number of keepalive packets a destination must receive to change a link's status from down to up. • **Down-count**—The number of keepalive packets a destination must fail to receive before the network takes down a link.	**extensive**
RE Keepalive statistics	Keepalive statistics for the packets handled by the Routing Engine. • **LCP echo req Tx**—LCP echo requests sent from the Routing Engine. • **LCP echo req Rx**—LCP echo requests received at the Routing Engine. • **LCP echo rep Tx**—LCP echo responses sent from the Routing Engine. • **LCP echo rep Rx**—LCP echo responses received at the Routing Engine. • **LCP echo req timeout**—Number of keepalive packets where the keepalive aging timer has expired. • **LCP Rx echo req Magic Num Failures**—LCP echo requests where the magic numbers shared between the PPP peers during LCP negotiation did not match. • **LCP Rx echo rep Magic Num Failures**—LCP echo responses where the magic numbers shared between the PPP peers during LCP negotiation did not match.	**extensive**

Table 24: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output

LCP	LCP information:
	• **State**—LCP protocol state (all platforms except M120 and M320 routers):
	  • **Ack-rcvd**—A Configure-Request has been sent and a Configure-Ack has been received.
	  • **Ack-sent**—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.
	  • **Closed**—Link is not available for traffic.
	  • **Opened**—Link is administratively available for traffic.
	  • **Req-sent**—An attempt has been made to configure the connection.
	• **State**—LCP protocol state (M120 and M320 routers):
	  • **Ack-rcvd**—A Configure-Request has been sent and a Configure-Ack has been received.
	  • **Ack-sent**—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.
	  • **Closed**—Link is available (up), but no Open has occurred.
	  • **Closing**—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.
	  • **Opened**—Link is administratively available for traffic. A Configure-Ack has been both sent and received.
	  • **Req-sent**—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received.
	  • **Starting**—An administrative Open has been initiated, but the lower layer is still unavailable (Down).
	  • **Stopped**—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack.
	  • **Stopping**—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.
	• **Last started**—LCP state start time.
	• **Last completed**—LCP state completion time.
		extensive

Table 24: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output

	• **Negotiated options**:
	  • **ACFC**—Address and-Control Field Compression. A configuration option that provides a method to negotiate the compression of the Data Link Layer Address and Control fields.
	  • **Asynchronous map**—Asynchronous control character map. A configuration option used on asynchronous links such as telephone lines to identify control characters that must be replaced by a two-character sequence to prevent them from being interpreted by equipment used to establish the link.
	  • **Authentication protocol**—Protocol used for authentication. This option provides a method to negotiate the use of a specific protocol for authentication. It requires a peer to authenticate itself before allowing network-layer protocol packets to be exchanged. By default, authentication is not required.
	  • **Authentication algorithm**—Type of authentication algorithm. The Message Digest algorithm (MD5) is the only algorithm supported.
	  • **Endpoint discriminator class**—For multilink PPP (MLPPP), a configuration option that identifies the system transmitting the packet. This option advises a system that the peer on this link could be the same as the peer on another existing link.
	  • **Magic number**—A configuration option that provides a method to detect looped-back links and other data-link layer anomalies. By default, the magic number is not negotiated.
	  • **MRU**—Maximum receive unit. A configuration option that may be sent to inform the peer that the implementation can receive larger packets, or to request that the peer send smaller packets. The default value is 1500 octets.
	  • **MRRU**—For multilink PPP, the maximum receive reconstructed unit. A configuration option that specifies the maximum number of octets in the Information fields of reassembled packets.
	  • **Multilink header suspendable classes**—For MLPPP, an LCP option that advises the peer that the implementation wishes to receive fragments with a format given by the code number, with the maximum number of suspendable classes given.
	  • **Multilink header format classes**—For MLPPP, an LCP option that advises the peer that the implementation wishes to receive fragments with a format given by the code number.
	  • **PFC**—Protocol-Field-Compression. A configuration option that provides a method to negotiate the compression of the PPP Protocol field.
	  • **short sequence**—For MLPPP, an option that advises the peer that the implementation wishes to receive fragments with short, 12-bit sequence numbers.	

Table 24: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
Authentication	CHAP or PAP authentication state information. For CHAP authentication: • **Chap-ans-rcvd**—Packet was sent from the peer, indicating that the peer received the **Chap-resp-sent** packet. • **Chap-ans-sent**—Packet was sent from the authenticator, indicating that the authenticator received the peer's **Chap-resp-rcvd** packet. • **Chap-chal-rcvd**—Challenge packet has been received by the peer. • **Chap-chal-sent**—Challenge packet has been sent by the authenticator to begin the CHAP protocol or has been transmitted at any time during the Network-Layer Protocol (NCP) phase to ensure that the connection has not been altered. • **Chap-resp-rcvd**—CHAP response packet has been received by the authenticator. • **Chap-resp-sent**—CHAP response packet has been sent to the authenticator. • **Closed**—Link is not available for authentication. • **Failure**—Authenticator compares the response value in the response packet from the peer with its own response value, but the value does not match. Authentication fails. • **Success**—Authenticator compares the response value in the response packet from the peer with its own response value, and the value matches. Authentication is successful. For PAP authentication: • **Pap-resp-sent**—PAP response sent to peer (ACK/NACK)t. • **Pap-req-rcvd**—PAP request packet received from peer. • **Pap-resp-rcvd**—PAP response received from the peer (ACK/NACK). • **Pap-req-sent**—PAP request packet sent to the peer. • **Closed**—Link is not available for authentication. • **Failure**—Authenticator compares the response value in the response packet from the peer with its own response value, but the value does not match. Authentication fails. • **Success**—Authenticator compares the response value in the response packet from the peer with its own response value, and the value matches. Authentication is successful.	None specified

Table 24: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output

IPCP	Internet Protocol Control Protocol (IPCP) information.
	• **State**—(All platforms except M120 and M320 routers) One of the following values:
	  • **Ack-rcvd**—A Configure-Request has been sent and a Configure-Ack has been received.
	  • **Ack-sent**—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.
	  • **Closed**—Link is not available for traffic.
	  • **Opened**—Link is administratively available for traffic.
	  • **Req-sent**—An attempt has been made to configure the connection.
	• **State**—(M120 and M320 routers) One of the following values:
	  • **Ack-rcvd**—A Configure-Request has been sent and a Configure-Ack has been received.
	  • **Ack-sent**—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.
	  • **Closed**—Link is available (up), but no Open has occurred.
	  • **Closing**—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.
	  • **Opened**—Link is administratively available for traffic. A Configure-Ack has been both sent and received.
	  • **Req-sent**—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received.
	  • **Starting**—An administrative Open has been initiated, but the lower layer is still unavailable (Down).
	  • **Stopped**—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack.
	  • **Stopping**—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.
	• **Last started**—IPCP state start time.
	• **Last completed**—IPCP state authentication completion time.
	• **Negotiated options:**
	  • **compression protocol**—Negotiate the use of a specific compression protocol. By default, compression is not enabled.
	  • **local address**—Desired local address of the sender of a Configure-Request. If all four octets are set to zero, the peer provides the IP address.
	  • **primary DNS server**—Negotiate with the remote peer to select the address of the primary DNS server to be used on the local end of the link.
	  • **primary WINS server**—Negotiate with the remote peer to select the address of the primary WINS server to be used on the local end of the link.
	  • **remote address**—IP address of the remote end of the link in dotted quad notation.
	  • **secondary DNS server**—Negotiate with the remote peer to select the address of the secondary DNS server to be used on the local end of the link.
	  • **secondary WINS server**—Negotiate with the remote peer to select the address of the secondary WINS server to be used on the local end of the link.
		extensive

Table 24: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output

IPV6CP	Internet Protocol version 6 Control Protocol (IPV6CP) information.
	• **State**—(All platforms except M120 and M320 routers) One of the following values:
	  • **Ack-rcvd**—A Configure-Request has been sent and a Configure-Ack has been received.
	  • **Ack-sent**—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.
	  • **Closed**—Link is not available for traffic.
	  • **Opened**—Link is administratively available for traffic.
	  • **Req-sent**—An attempt has been made to configure the connection.
	• **State**—(M120 and M320 routers) One of the following values:
	  • **Ack-rcvd**—A Configure-Request has been sent and a Configure-Ack has been received.
	  • **Ack-sent**—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.
	  • **Closed**—Link is available (up), but no Open has occurred.
	  • **Closing**—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.
	  • **Opened**—Link is administratively available for traffic. A Configure-Ack has been both sent and received.
	  • **Req-sent**—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received.
	  • **Starting**—An administrative Open has been initiated, but the lower layer is still unavailable (Down).
	  • **Stopped**—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack.
	  • **Stopping**—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.
	• **Last started**—IPV6CP state start time.
	• **Last completed**—IPV6CP state authentication completion time.
	• **Negotiated options**:
	  • **local interface identifier**—Desired local address of the sender of a Configure-Request. If all four octets are set to zero, the peer provides the IP address.
	  • **remote interface identifier**—IP address of the remote end of the link in dotted quad notation.
		extensive

OSINLCP State	OSI Network Layer Control Protocol (OSINLCP) protocol state information (all platforms except M120 and M320 routers):
	• **State:**
	  • **Ack-rcvd**—Configure-Request has been sent and Configure-Ack has been received.
	  • **Ack-sent**—Configure-Request and Configure-Ack have both been sent, but Configure-Ack has not yet been received.
	  • **Closed**—Link is not available for traffic.
	  • **Opened**—Link is administratively available for traffic.
	  • **Req-sent**—Attempt has been made to configure the connection.
	• **Last started**—OSINLCP state start time.
	• **Last completed**—OSINCLP state completion time.
		extensive

Table 24: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output

TAGCP	TAGCP information.
	• **State**—(All platforms except M120 and M320 routers) One of the following values:
	  • **Ack-rcvd**—A Configure-Request has been sent and a Configure-Ack has been received.
	  • **Ack-sent**—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.
	  • **Closed**—Link is not available for traffic.
	  • **Opened**—Link is administratively available for traffic.
	  • **Req-sent**—An attempt has been made to configure the connection.
	• **State**—(M120 and M320 routers) One of the following values:
	  • **Ack-rcvd**—A Configure-Request has been sent and a Configure-Ack has been received.
	  • **Ack-sent**—A Configure-Request and a Configure-Ack have both been sent, but a Configure-Ack has not yet been received.
	  • **Closed**—Link is available (up), but no Open has occurred.
	  • **Closing**—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.
	  • **Opened**—Link is administratively available for traffic. A Configure-Ack has been both sent and received.
	  • **Req-sent**—An attempt has been made to configure the connection. A Configure-Request has been sent but a Configure-Ack has not yet been received.
	  • **Starting**—An administrative Open has been initiated, but the lower layer is still unavailable (Down).
	  • **Stopped**—The system is waiting for a Down event after the This-Layer-Finished action, or after sending a Terminate-Ack.
	  • **Stopping**—A Terminate-Request has been sent but a Terminate-Ack has not yet been received.
	• **Last started**—TAGCP state start time.
	• **Last completed**—TAGCP state authentication completion time.
		extensive
		none

Sample Output

show ppp interface

```
user@host> show ppp interface si-1/3/0.0
Session si-1/3/0.0, Type: PPP, Phase: Authenticate
  Session flags: Monitored
    LCP State: Opened
    AUTHENTICATION: CHAP State: Chap-resp-sent, Chap-ans-sent
    IPCP State: Closed, OSINLCP State: Closed
```

show ppp interface extensive

```
user@host> show ppp interface si-0/0/3.0 extensive
    Session si-0/0/3.0, Type: PPP, Phase: Network
        Keepalive settings: Interval 30 seconds, Up-count 1, Down-count 3
RE Keepalive statistics:
  LCP echo req Tx        :  657 (last sent 00:50:10 ago)
  LCP echo req Rx        :    0 (last seen: never)
  LCP echo rep Tx        :    0
```

```
  LCP echo rep Rx        :  657
  LCP echo req timout    :    0
  LCP Rx echo req Magic Num Failures :    0
  LCP Rx echo rep Magic Num Failures :    0
LCP
      State: Opened
      Last started: 2007-01-29 10:43:50 PST
      Last completed: 2007-01-29 10:43:50 PST
      Negotiated options:
        Authentication protocol: PAP, Magic number: 2341124815, MRU: 4470
    Authentication: PAP
      State: Success
      Last started: 2007-01-29 10:43:50 PST
      Last completed: 2007-01-29 10:43:50 PST
    IPCP
      State: Opened
      Last started: 2007-01-29 10:43:50 PST
      Last completed: 2007-01-29 10:43:50 PST
      Negotiated options:
        Local address: 10.10.10.1, Remote address: 10.10.10.2
```

show ppp interface terse

```
user@host> show ppp interface si-1/3/0 terse
Session name    Session type   Session phase      Session flags
si-1/3/0.0      PPP            Authenticate       Monitored
```

CHAPTER 11

Subscriber Management Dynamic Protocol CLI Commands

show igmp interface

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.

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*

List of Sample Output

show igmp interface on page 312
show igmp interface brief on page 312
show igmp interface detail on page 313
show igmp interface <interface-name> on page 313

Output Fields

Table 25 on page 310 describes the output fields for the **show igmp interface** command. Output fields are listed in the approximate order in which they appear.

Table 25: 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

Table 25: show igmp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
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

Table 25: show igmp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
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: 10.111.30.1
    State:         Up Timeout:    None Version:  2 Groups:      4
    SSM Map Policy: ssm-policy-A
Interface: so-1/0/0.0
    Querier: 10.111.10.1
    State:         Up Timeout:    None Version:  2 Groups:      2
    SSM Map Policy: ssm-policy-B
Interface: so-1/0/1.0
    Querier: 10.111.20.1
    State:         Up Timeout:    None Version:  2 Groups:      4
    SSM Map Policy: ssm-policy-C
Immediate Leave: On
Promiscuous Mode: Off

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 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 312.

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 312.

show igmp interface <interface-name>

```
user@host# show igmp interface ge-3/2/0.0
Interface: ge-3/2/0.0
  Querier: 20.1.1.1
  State: Up Timeout:    None Version:  3 Groups:      1
  Group limit: 8
  Group threshold: 60
  Group log-interval: 10
  Immediate leave:  Off
  Promiscuous mode:  Off
```

show igmp statistics

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.

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*

List of Sample Output

show igmp statistics on page 315
show igmp statistics interface on page 316

Output Fields

Table 26 on page 314 describes the output fields for the **show igmp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 26: 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.

Table 26: 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).
Received	Number of messages received.
Sent	Number of messages sent.
Rx errors	Number of received packets that contained errors.
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.

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
```

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
```

CHAPTER 12

Subscriber Management Subscriber CLI Commands

show subscribers

Syntax

show subscribers
<detail | extensive | terse>
<aci-interface-set-name *aci-interface-set-name*>
<address *address*>
<agent-circuit-identifier *agent-circuit-identifier-substring*>
<client-type *client-type*>
<count>
<interface *interface*>
<logical-system *logical-system*>
<mac-address *mac-address*>
<physical-interface *physical-interface-name*>
<profile-name *profile-name*>
<routing-instance *routing-instance*>
<stacked-vlan-id *stacked-vlan-id*>
<subscriber-state *subscriber-state*>
<user-name *user-name*>
<vci *vci-identifier*>
<vpi *vpi-identifier*>
<vlan-id *vlan-id*>

Release Information

Command introduced in Junos OS Release 9.3.
Command introduced in Junos OS Release 9.3 for EX Series switches.
client-type, **mac-address**, **subscriber-state**, and **extensive** options introduced in Junos OS Release 10.2.
count option usage with other options introduced in Junos OS Release 10.2.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Options **aci-interface-set-name** and **agent-circuit-identifier** introduced in Junos OS Release 12.2.
The **physical-interface** and **user-name** options introduced in Junos OS Release 12.3.
Options **vci** and **vpi** introduced in Junos OS Release 12.3R3 and supported in later 12.3R*x* releases.
Options **vci** and **vpi** supported in Junos OS Release 13.2 and later releases. (Not supported in Junos OS Release 13.1.)

Description

Display information for active subscribers.

Options

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

aci-interface-set-name—(Optional) Display all dynamic subscriber sessions that use the specified agent circuit identifier (ACI) interface set. Use the ACI interface set name generated by the router, such as aci-1003-ge-1/0/0.4001, and not the actual ACI value found in the DHCP or PPPoE control packets.

address—(Optional) Display subscribers whose IP address matches the specified address. You must specify the IPv4 or IPv6 address prefix without a netmask (for example, 192.168.17.1). If you specify the IP address as a prefix with a netmask (for example, 192.168.17.1/32), the router displays a message that the IP address is invalid, and rejects the command.

agent-circuit-identifier-substring—(Optional) Display all dynamic subscriber sessions whose ACI value matches the specified substring.

client-type—(Optional) Display subscribers whose client type matches the specified client type (DHCP, L2TP, PPP, PPPOE, VLAN, or static).

count—(Optional) Display the count of total subscribers and active subscribers for any specified option. You can use the **count** option alone or with the **address**, **client-type**, **interface**, **logical-system**, **mac-address**, **profile-name**, **routing-instance**, **stacked-vlan-id**, **subscriber-state**, or **vlan-id** options.

id—(Optional) Display a specific subscriber session whose session id matches the specified subscriber ID. You can display subscriber IDs by using the **show subscribers extensive** or the **show subscribers interface extensive** commands.

interface—(Optional) Display subscribers whose interface matches the specified interface.

logical-system—(Optional) Display subscribers whose logical system matches the specified logical system.

mac-address—(Optional) Display subscribers whose MAC address matches the specified MAC address.

physical-interface-name—(M120, M320, and MX Series routers only) (Optional) Display subscribers whose physical interface matches the specified physical interface.

profile-name—(Optional) Display subscribers whose dynamic profile matches the specified profile name.

routing-instance—(Optional) Display subscribers whose routing instance matches the specified routing instance.

subscriber-state—(Optional) Display subscribers whose subscriber state matches the specified subscriber state (ACTIVE, CONFIGURED, INIT, TERMINATED, or TERMINATING).

user-name—(M120, M320, and MX Series routers only) (Optional) Display subscribers whose username matches the specified subscriber name.

vci-identifier—(MX Series routers with MPCs and ATM MICs with SFP only) (Optional) Display active ATM subscribers whose ATM virtual circuit identifier (VCI) matches the specified VCI identifier. The range of values is **0** through **255**.

vpi-identifier—(MX Series routers with MPCs and ATM MICs with SFP only) (Optional) Display active ATM subscribers whose ATM virtual path identifier (VPI) matches the specified VPI identifier. The range of values is **0** through **65535**.

vlan-id—(Optional) Display subscribers whose VLAN ID matches the specified VLAN ID.

stacked-vlan-id—(Optional) Display subscribers whose stacked VLAN ID matches the specified stacked VLAN ID.

NOTE: Due to display limitations, logical system and routing instance output values are truncated when necessary.

Required Privilege Level view

Related Documentation

- show subscribers summary on page 336
- *Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*

List of Sample Output

show subscribers vlan-id detail on page 335
show subscribers vpi vci extensive (PPPoE-over-ATM Subscriber Session) on page 335

Output Fields Table 27 on page 321 lists the output fields for the **show subscribers** command. Output fields are listed in the approximate order in which they appear.

Table 27: show subscribers Output Fields

Field Name	Field Description
Interface	Interface associated with the subscriber. The router or switch displays subscribers whose interface matches or begins with the specified interface. The * character indicates a continuation of addresses for the same session.
IP Address/VLAN ID	Subscriber IP address or VLAN ID associated with the subscriber in the form *tpid.vlan-id* No IP address or VLAN ID is assigned to an L2TP tunnel-switched session. For these subscriber sessions the value is **Tunnel-switched**.
User Name	Name of subscriber.
LS:RI	Logical system and routing instance associated with the subscriber.
Type	Subscriber client type (DHCP, L2TP, PPP, PPPoE, STATIC-INTERFACE, VLAN).
IP Address	Subscriber IPv4 address.
IP Netmask	Subscriber IP netmask.
Primary DNS Address	IP address of primary DNS server.
Secondary DNS Address	IP address of secondary DNS server.
Primary WINS Address	IP address of primary WINS server.
Secondary WINS Address	IP address of secondary WINS server.
IPv6 Address	Subscriber IPv6 address, or multiple addresses.
IPv6 Prefix	Subscriber IPv6 prefix. If you are using DHCPv6 prefix delegation, this is the delegated prefix.
IPv6 User Prefix	IPv6 prefix obtained through ND/RA.
IPv6 Address Pool	Subscriber IPv6 address pool. The IPv6 address pool is used to allocate IPv6 prefixes to the DHCPv6 clients.
IPv6 Network Prefix Length	Length of the network portion of the IPv6 address.
IPv6 Prefix Length	Length of the subscriber IPv6 prefix.

Table 27: show subscribers Output Fields *(continued)*

Field Name	Field Description
Logical System	Logical system associated with the subscriber.
Routing Instance	Routing instance associated with the subscriber.
Interface Type	Whether the subscriber interface is **Static** or **Dynamic**.
Interface Set	Internally generated name of the dynamic ACI interface set used by the subscriber session.
Interface Set Type	Interface type of the ACI interface set: **Dynamic**. This is the only ACI interface set type currently supported.
Interface Set Session ID	Identifier of the dynamic ACI interface set entry in the session database.
Underlying Interface	Name of the underlying interface for the subscriber session.
Dynamic Profile Name	Dynamic profile used for the subscriber.
Dynamic Profile Version	Version number of the dynamic profile used for the subscriber.
MAC Address	MAC address associated with the subscriber.
State	Current state of the subscriber session (**Init**, **Configured**, **Active**, **Terminating**, **Tunneled**).
L2TP State	Current state of the L2TP session, **Tunneled** or **Tunnel-switched**. When the value is **Tunnel-switched**, two entries are displayed for the subscriber; the first entry is at the LNS interface on the LTS and the second entry is at the LAC interface on the LTS.
Tunnel switch Profile Name	Name of the L2TP tunnel switch profile that initiates tunnel switching.
Local IP Address	IP address of the local gateway (LAC).
Remote IP Address	IP address of the remote peer (LNS).
VLAN Id	VLAN ID associated with the subscriber in the form *tpid.vlan-id*.
Stacked VLAN Id	Stacked VLAN ID associated with the subscriber in the form *tpid.vlan-id*.
RADIUS Accounting ID	RADIUS accounting ID associated with the subscriber.
Agent Circuit ID	Option 82 agent circuit ID associated with the subscriber. The ID is displayed as an ASCII string unless the value has nonprintable characters, in which case it is displayed in hexadecimal format.
Agent Remote ID	Option 82 agent remote ID associated with the subscriber. The ID is displayed as an ASCII string unless the value has nonprintable characters, in which case it is displayed in hexadecimal format.
DHCP Relay IP Address	IP address used by the DHCP relay agent.

Table 27: show subscribers Output Fields *(continued)*

Field Name	Field Description
ATM VPI	(MX Series routers with MPCs and ATM MICs with SFP only) ATM virtual path identifier (VPI) on the subscriber's physical interface.
ATM VCI	(MX Series routers with MPCs and ATM MICs with SFP only) ATM virtual circuit identifier (VCI) for each VPI configured on the subscriber interface.
Login Time	Date and time at which the subscriber logged in.
Effective shaping-rate	Actual downstream traffic shaping rate for the subscriber, in kilobits per second.
IPv4 rpf-check Fail Filter Name	Name of the filter applied by the dynamic profile to IPv4 packets that fail the RPF check.
IPv6 rpf-check Fail Filter Name	Name of the filter applied by the dynamic profile to IPv6 packets that fail the RPF check.
DHCP Options	**len** = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCP options, as defined in RFC 2132.
Session ID	ID number for a subscriber service session.
Underlying Session ID	For DHCPv6 subscribers on a PPPoE network, displays the session ID of the underlying PPPoE interface.
Service Sessions	Number of service sessions (that is, a service activated using RADIUS CoA) associated with the subscribers.
Service Session Name	Service session profile name.
Session Timeout (seconds)	Number of seconds of access provided to the subscriber before the session is automatically terminated.
Idle Timeout (seconds)	Number of seconds subscriber can be idle before the session is automatically terminated.
IPv6 Delegated Address Pool	Name of the pool used for DHCPv6 prefix delegation.
IPv6 Delegated Network Prefix Length	Length of the prefix configured for the IPv6 delegated address pool.
IPv6 Interface Address	Address assigned by the Framed-Ipv6-Prefix AAA attribute.
IPv6 Framed Interface Id	Interface ID assigned by the Framed-Interface-Id AAA attribute.
ADF IPv4 Input Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv4 input filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.

Table 27: show subscribers Output Fields *(continued)*

Field Name	Field Description
ADF IPv4 Output Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv4 output filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv6 Input Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv6 input filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
ADF IPv6 Output Filter Name	Name assigned to the Ascend-Data-Filter (ADF) interface IPv6 output filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
IPv4 Input Filter Name	Name assigned to the IPv4 input filter (client or service session).
IPv4 Output Filter Name	Name assigned to the IPv4 output filter (client or service session).
IPv6 Input Filter Name	Name assigned to the IPv6 input filter (client or service session).
IPv6 Output Filter Name	Name assigned to the IPv6 output filter (client or service session).
IFL Input Filter Name	Name assigned to the logical interface input filter (client or service session).
IFL Output Filter Name	Name assigned to the logical interface output filter (client or service session).

Sample Output

show subscribers (IPv4)

```
user@host> show subscribers
Interface                IP Address/VLAN ID   User Name         LS:RI
ge-1/3/0.1073741824      100                                    default:default
demux0.1073741824        100.0.0.10           WHOLESALER-CLIENT default:default
demux0.1073741825        101.0.0.3            RETAILER1-CLIENT  test1:retailer1
demux0.1073741826        102.0.0.3            RETAILER2-CLIENT  test1:retailer2
```

show subscribers (IPv6)

```
user@host> show subscribers
Interface           IP Address/VLAN ID   User Name           LS:RI
ge-1/0/0.0          2001::c0:0:0:0/74    WHOLESALER-CLIENT   default:default
     *              2002::1/128          subscriber-25       default:default
```

show subscribers (IPv4 and IPv6 Dual Stack)

```
user@host> show subscribers
Interface            IP Address/VLAN ID         User Name
LS:RI
demux0.1073741834    0x8100.1002 0x8100.1
default:default
demux0.1073741835    0x8100.1001 0x8100.1
default:default
pp0.1073741836       61.1.1.1                   dualstackuser1@ISP1.com
```

```
default:ASP-1
*                   2041:1:1::/48
*                   2061:1:1:1::/64
pp0.1073741837      23.1.1.3                dualstackuser2@ISP1.com
default:ASP-1
*                   2001:1:2:5::/64
```

show subscribers (LNS on MX Series Routers)

```
user@host> show subscribers
Interface          IP Address/VLAN ID   User Name         LS:RI
si-4/0/0.1         192.168.4.1          xyz@example.com   default:default
```

show subscribers (L2TP Switched Tunnels)

```
user@host> show subscribers
Interface           IP Address/VLAN ID    User Name              LS:RI
si-2/1/0.1073741842 Tunnel-switched       ap@lts.com        default:default

si-2/1/0.1073741843 Tunnel-switched       ap@lts.com        default:default
```

show subscribers client-type dhcp detail

```
user@host> show subscribers client-type dhcp detail
Type: DHCP
IP Address: 100.20.9.7
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:10:95:00:00:98
State: Active
Radius Accounting ID: jnpr :2304
Login Time: 2009-08-25 14:43:52 PDT

Type: DHCP
IP Address: 100.20.10.7
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744383
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:10:94:00:01:f3
State: Active
Radius Accounting ID: jnpr :2560
Login Time: 2009-08-25 14:43:56 PDT
```

show subscribers count

```
user@host> show subscribers count
Total Subscribers: 188, Active Subscribers: 188
```

show subscribers address detail (IPv6)

```
user@host> show subscribers address 100.16.12.137 detail
```

```
Type: PPPoE
User Name: pppoeTerV6User1Svc
IP Address: 100.16.12.137
IP Netmask: 255.0.0.0
IPv6 User Prefix: 1016:0:0:c88::/64
Logical System: default
Routing Instance: default
Interface: pp0.1073745151
Interface type: Dynamic
Underlying Interface: demux0.8201
Dynamic Profile Name: pppoe-client-profile
MAC Address: 00:0d:02:01:00:01
Session Timeout (seconds): 31622400
Idle Timeout (seconds): 86400
State: Active
Radius Accounting ID: jnpr demux0.8201:6544
Session ID: 6544
Agent Circuit ID: ifl3720
Agent Remote ID: ifl3720
Login Time: 2012-05-21 13:37:27 PDT
Service Sessions: 1
```

show subscribers detail (IPv4)

```
user@host> show subscribers detail
Type: DHCP
IP Address: 100.20.9.7
IP Netmask: 255.255.0.0
Primary DNS Address: 192.168.17.1
Secondary DNS Address: 192.168.17.2
Primary WINS Address: 192.168.22.1
Secondary WINS Address: 192.168.22.2
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:10:95:00:00:98
State: Active
Radius Accounting ID: jnpr :2304
Session Timeout (seconds): 3600
Idle Timeout (seconds): 600
Login Time: 2009-08-25 14:43:52 PDT
DHCP Options: len 52
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 08 33 04 00 00
00 3c 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 36 2f
33 2d 37 2d 30 37 05 01 06 0f 21 2c
Service Sessions: 2
```

show subscribers detail (IPv6)

```
user@host> show subscribers detail
Type: DHCP
User Name: pd-user1
IPv6 Prefix: 2002:db2:ffff:1::/64
Logical System: default
Routing Instance: default
Interface: ge-3/1/3.2
Interface type: Static
MAC Address: 00:51:ff:ff:00:03
State: Active
```

```
Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-08-25 12:12:26 PDT
DHCP Options: len 42
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 51 ff ff 00 03
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00 00
00 00
```

show subscribers detail (IPv6 Static Demux Interface)

```
user@host> show subscribers detail
Type: STATIC-INTERFACE
User Name: demux0.1@jnpr.net
IPv6 Prefix: 1:2:3:4:5:6:7:aa/128
Logical System: default
Routing Instance: default
Interface: demux0.1
Interface type: Static
Dynamic Profile Name: junos-default-profile
State: Active
Radius Accounting ID: 185
Login Time: 2010-05-18 14:33:56 EDT
```

show subscribers detail (L2TP LNS Subscribers on MX Series Routers)

```
user@host> show subscribers detail
Type: L2TP
User Name: user1@jnpr.net
IP Address: 10.1.32.58
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: si-5/2/0.1073749824
Interface type: Dynamic
Dynamic Profile Name: dyn-lns-profile2
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 8001
Session ID: 8001
Login Time: 2011-04-25 20:27:50 IST
```

show subscribers detail (L2TP Switched Tunnels)

```
user@host> show subscribers detail
Type: L2TP
User Name: ap@example.com
Logical System: default
Routing Instance: default
Interface: si-2/1/0.1073741842
Interface type: Dynamic
Dynamic Profile Name: dyn-lts-profile
State: Active
L2TP State: Tunnel-switched
Tunnel switch Profile Name: ce-lts-profile
Local IP Address: 10.50.1.1
Remote IP Address: 192.168.20.3
Radius Accounting ID: 21
Session ID: 21
Login Time: 2013-01-18 03:01:11 PST

Type: L2TP
User Name: ap@example.com
```

```
Logical System: default
Routing Instance: default
Interface: si-2/1/0.1073741843
Interface type: Dynamic
Dynamic Profile Name: dyn-lts-profile
State: Active
L2TP State: Tunnel-switched
Tunnel switch Profile Name: ce-lts-profile
Local IP Address: 10.30.1.1
Remote IP Address: 172.20.1.10
Session ID: 22
Login Time: 2013-01-18 03:01:14 PST
```

show subscribers detail (Tunneled Subscriber)

```
user@host> show subscribers detail
Type: PPPoE
User Name: user1@example.com
Logical System: default
Routing Instance: default
Interface: pp0.1
State: Active, Tunneled
Radius Accounting ID: 512
```

show subscribers detail (IPv4 and IPv6 Dual Stack)

```
user@host> show subscribers detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlanProfile
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.1001
VLAN Id: 0x8100.1
Login Time: 2011-11-30 00:18:04 PST

Type: PPPoE
User Name: dualstackuser1@ISP1.com
IP Address: 61.1.1.1
IPv6 Prefix: 2041:1:1::/48
IPv6 User Prefix: 2061:1:1:1::/64
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Dynamic
Dynamic Profile Name: dualStack-Profile1
MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: 2
Session ID: 2
Login Time: 2011-11-30 00:18:05 PST

Type: DHCP
IPv6 Prefix: 2041:1:1::/48
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Static
```

```
MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: jnpr :3
Session ID: 3
Underlying Session ID: 2
Login Time: 2011-11-30 00:18:35 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 00 64 03 01 02
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00 00
00 00
```

show subscribers detail (ACI Interface Set Session)

```
user@host> show subscribers detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0
Interface Set: aci-1001-ge-1/0/0.2800
Interface Set Session ID: 0
Underlying Interface: ge-1/0/0.2800
Dynamic Profile Name: aci-vlan-set-profile-2
Dynamic Profile Version: 1
State: Active
Session ID: 1
Agent Circuit ID: aci-ppp-dhcp-20
Login Time: 2012-05-26 01:54:08 PDT
```

show subscribers detail (PPPoE Subscriber Session with ACI Interface Set)

```
user@host> show subscribers detail
Type: PPPoE
User Name: ppphint2
IP Address: 10.10.1.5
Logical System: default
Routing Instance: default
Interface: pp0.1073741825
Interface type: Dynamic
Interface Set: aci-1001-demux0.1073741824
Interface Set Type: Dynamic
Interface Set Session ID: 2
Underlying Interface: demux0.1073741824
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address: 00:00:64:39:01:02
State: Active
Radius Accounting ID: 3
Session ID: 3
Agent Circuit ID: aci-ppp-dhcp-dvlan-50
Login Time: 2012-03-07 13:46:53 PST
```

show subscribers extensive

```
user@host> show subscribers extensive
Type: DHCP
User Name: pd-user1
IPv6 Prefix: 2002:db2:ffff:1::/64
Logical System: default
Routing Instance: default
Interface: ge-3/1/3.2
Interface type: Static
```

```
MAC Address: 00:51:ff:ff:00:03
State: Active
Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-08-25 12:12:26 PDT
DHCP Options: len 42
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 51 ff ff 00 03
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00 00
00 00
IPv6 Address Pool: pd_pool
IPv6 Network Prefix Length: 48
```

show subscribers extensive (RPF Check Fail Filter)

```
user@host> show subscribers extensive
...
Type: VLAN
  Logical System: default
  Routing Instance: default
  Interface: ae0.1073741824
  Interface type: Dynamic
  Dynamic Profile Name: vlan-prof
  State: Active
  Session ID: 9
  VLAN Id: 100
  Login Time: 2011-08-26 08:17:00 PDT
  IPv4 rpf-check Fail Filter Name: rpf-allow-dhcp
  IPv6 rpf-check Fail Filter Name: rpf-allow-dhcpv6
...
```

show subscribers extensive (L2TP LNS Subscribers on MX Series Routers)

```
user@host> show subscribers extensive
Type: L2TP
User Name: user1@jnpr.net
IP Address: 10.1.32.58
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: si-5/2/0.1073749824
Interface type: Dynamic
Dynamic Profile Name: dyn-lns-profile2
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 8001
Session ID: 8001
Login Time: 2011-04-25 20:27:50 IST
IPv4 Input Filter Name: classify-si-5/2/0.1073749824-in
IPv4 Output Filter Name: classify-si-5/2/0.1073749824-out
```

show subscribers extensive (IPv4 and IPv6 Dual Stack)

```
user@host> show subscribers extensive
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlanProfile
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.1001
```

```
VLAN Id: 0x8100.1
Login Time: 2011-11-30 00:18:04 PST

Type: PPPoE
User Name: dualstackuser1@ISP1.com
IP Address: 61.1.1.1
IPv6 Prefix: 2041:1:1::/48
IPv6 User Prefix: 2061:1:1:1::/64
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Dynamic
Dynamic Profile Name: dualStack-Profile1
MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: 2
Session ID: 2
Login Time: 2011-11-30 00:18:05 PST
IPv6 Delegated Network Prefix Length: 48
IPv6 Interface Address: 2061:1:1:1::1/64
IPv6 Framed Interface Id: 1:1:2:2
IPv4 Input Filter Name: FILTER-IN-pp0.1073741825-in
IPv4 Output Filter Name: FILTER-OUT-pp0.1073741825-out
IPv6 Input Filter Name: FILTER-IN6-pp0.1073741825-in
IPv6 Output Filter Name: FILTER-OUT6-pp0.1073741825-out

Type: DHCP
IPv6 Prefix: 2041:1:1::/48
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Static
MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: jnpr :3
Session ID: 3
Underlying Session ID: 2
Login Time: 2011-11-30 00:18:35 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 00 64 03 01 02
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00 00
00 00
IPv6 Delegated Network Prefix Length: 48
```

show subscribers extensive (Effective Shaping-Rate)

```
user@host> show subscribers extensive
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741837
Interface type: Dynamic
Interface Set: ifset-1
Underlying Interface: ae1
Dynamic Profile Name: svlan-dhcp-test
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.201
VLAN Id: 0x8100.201
Login Time: 2011-11-30 00:18:04 PST
```

```
Effective shaping-rate: 31000000k
...
```

show subscribers aci-interface-set-name detail (Subscriber Sessions Using Specified ACI Interface Set)

```
user@host> show subscribers aci-interface-set-name aci-1003-ge-1/0/0.4001 detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-set-profile
Dynamic Profile Version: 1
State: Active
Session ID: 13
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:56 PDT

Type: PPPoE
User Name: ppphint2
IP Address: 10.10.1.7
Logical System: default
Routing Instance: default
Interface: pp0.1073741834
Interface type: Dynamic
Interface Set: aci-1003-ge-1/0/0.4001
Interface Set Type: Dynamic
Interface Set Session ID: 13
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address: 00:00:65:26:01:02
State: Active
Radius Accounting ID: 14
Session ID: 14
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:57 PDT
```

show subscribers agent-circuit-identifier detail (Subscriber Sessions Using Specified ACI Substring)

```
user@host> show subscribers agent-circuit-identifier aci-ppp-vlan detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-set-profile
Dynamic Profile Version: 1
State: Active
Session ID: 13
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:56 PDT

Type: PPPoE
User Name: ppphint2
IP Address: 10.10.1.7
Logical System: default
Routing Instance: default
Interface: pp0.1073741834
Interface type: Dynamic
Interface Set: aci-1003-ge-1/0/0.4001
```



```
Interface Set Type: Dynamic
Interface Set Session ID: 13
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address: 00:00:65:26:01:02
State: Active
Radius Accounting ID: 14
Session ID: 14
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:57 PDT
```

show subscribers interface extensive

```
user@host> show subscribers interface demux0.1073741826 extensive
Type: VLAN
User Name: test1@test.com
Logical System: default
Routing Instance: testnet
Interface: demux0.1073741826
Interface type: Dynamic
Dynamic Profile Name: profile-vdemux-relay-23qos
MAC Address: 00:00:6e:56:01:04
State: Active
Radius Accounting ID: 12
Session ID: 12
Stacked VLAN Id: 0x8100.1500
VLAN Id: 0x8100.2902
Login Time: 2011-10-20 16:21:59 EST

Type: DHCP
User Name: test1@test.com
IP Address: 172.16.200.6
IP Netmask: 255.255.255.0
Logical System: default
Routing Instance: testnet
Interface: demux0.1073741826
Interface type: Static
MAC Address: 00:00:6e:56:01:04
State: Active
Radius Accounting ID: 21
Session ID: 21
Login Time: 2011-10-20 16:24:33 EST
Service Sessions: 2

Service Session ID: 25
Service Session Name: SUB-QOS
State: Active

Service Session ID: 26
Service Session Name: service-cb-content
State: Active
IPv4 Input Filter Name: content-cb-in-demux0.1073741826-in
IPv4 Output Filter Name: content-cb-out-demux0.1073741826-out
```

show subscribers logical-system terse

```
user@host> show subscribers logical-system test1 terse
```

```
Interface           IP Address/VLAN ID   User Name         LS:RI
demux0.1073741825     101.0.0.3            RETAILER1-CLIENT  test1:retailer1
demux0.1073741826     102.0.0.3            RETAILER2-CLIENT  test1:retailer2
```

show subscribers physical-interface count

```
user@host> show subscribers physical-interface ge-1/0/0 count
Total subscribers: 3998, Active Subscribers: 3998
```

show subscribers routing-instance inst1 count

```
user@host> show subscribers routing-instance inst1 count
Total Subscribers: 188, Active Subscribers: 183
```

show subscribers stacked-vlan-id detail

```
user@host> show subscribers stacked-vlan-id 101 detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

show subscribers stacked-vlan-id vlan-id detail (Combined Output)

```
user@host> show subscribers stacked-vlan-id 101 vlan-id 100 detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

show subscribers stacked-vlan-id vlan-id interface detail (Combined Output for a Specific Interface)

```
user@host> show subscribers stacked-vlan-id 101 vlan-id 100 interface ge-1/2/0.* detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

show subscribers user-name detail

```
user@host> show subscribers user-name larry1 detail
Type: DHCP
User Name: larry1
IP Address: 100.0.0.37
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.1
Interface type: Static
Dynamic Profile Name: foo
```

```
MAC Address: 00:10:94:00:00:01
State: Active
Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-11-07 08:25:59 PST
DHCP Options: len 52
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 01 33 04 00 00
00 3c 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 32 2f
37 2d 30 2d 30 37 05 01 06 0f 21 2c
```

show subscribers vlan-id

```
user@host> show subscribers vlan-id 100
Interface           IP Address                              User Name
ge-1/0/0.1073741824
ge-1/2/0.1073741825
```

show subscribers vlan-id detail

```
user@host> show subscribers vlan-id 100 detail
Type: VLAN
Interface: ge-1/0/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: vlan-prof-tpid
State: Active
VLAN Id: 100
Login Time: 2009-03-11 06:48:54 PDT

Type: VLAN
Interface: ge-1/2/0.1073741825
Interface type: Dynamic
Dynamic Profile Name: vlan-prof-tpid
State: Active
VLAN Id: 100
Login Time: 2009-03-11 06:48:54 PDT
```

show subscribers vpi vci extensive (PPPoE-over-ATM Subscriber Session)

```
user@host> show subscribers vpi 40 vci 50 extensive
Type: PPPoE
User Name: testuser
IP Address: 100.0.0.2
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: pp0.0
Interface type: Static
MAC Address: 00:00:65:23:01:02
State: Active
Radius Accounting ID: 2
Session ID: 2
ATM VPI: 40
ATM VCI: 50
Login Time: 2012-12-03 07:49:26 PST
IP Address Pool: pool_1
IPv6 Framed Interface Id: 200:65ff:fe23:102
```

show subscribers summary

Syntax

show subscribers summary
< detail | extensive | terse>
<count>
physical-interface *physical-interface-name*
<all | logical-system *logical-system* pic | port | routing-instance *routing-instance* |slot>

Release Information

Command introduced in Junos OS Release 10.2.

Description

Display summary information for subscribers.

Options

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

count—(Optional) Display the count of total subscribers and active subscribers for any specified option.

logical-system—(Optional) Display subscribers whose logical system matches the specified logical system.

physical-interface-name—(M120, M320, and MX Series routers only) (Optional) Display a count of subscribers whose physical interface matches the specified physical interface, by subscriber state, client type and LS:RI.

pic—(M120, M320, and MX Series routers only) (Optional) Display a count of subscribers by PIC number and the total number of subscribers.

port—(M120, M320, and MX Series routers only) (Optional) Display a count of subscribers by port number and the total number of subscribers.

routing-instance—(Optional) Display subscribers whose routing instance matches the specified routing instance.

slot—(M120, M320, and MX Series routers only) (Optional) Display a count of subscribers by FPC slot number and the total number of subscribers.

NOTE: **Due to display limitations, logical system and routing instance output values are truncated when necessary.**

Required Privilege Level

view

Related Documentation

- show subscribers on page 318

List of Sample Output

show subscribers summary on page 338
show subscribers summary all on page 338
show subscribers summary physical-interface on page 338
show subscribers summary physical-interface pic on page 339

show subscribers summary physical-interface port on page 339
show subscribers summary physical-interface slot on page 339
show subscribers summary pic on page 339
show subscribers summary pic (Aggregated Ethernet Interfaces) on page 340
show subscribers summary port on page 340
show subscribers summary slot on page 340
show subscribers summary terse on page 340

Output Fields Table 28 on page 337 lists the output fields for the **show subscribers** command. Output fields are listed in the approximate order in which they appear.

Table 28: show subscribers Output Fields

Field Name	Field Description
Subscribers by State	Number of subscribers summarized by state. The summary information includes the following: • Init—Number of subscriber currently in the initialization state. • Configured—Number of configured subscribers. • Active—Number of active subscribers. • Terminating—Number of subscribers currently terminating. • Terminated—Number of terminated subscribers. • Total—Total number of subscribers for all states.
Subscribers by Client Type	Number of subscribers summarized by client type. Client types can include DHCP, L2TP, PPP, PPPOE, STATIC-INTERFACE, and VLAN. Also displays the total number of subscribers for all client types (Total).
Subscribers by LS:RI	Number of subscribers summarized by logical system:routing instance (LS:RI) combination. Also displays the total number of subscribers for all LS:RI combinations (Total).
Interface	Interface associated with the subscriber. The router or switch displays subscribers whose interface matches or begins with the specified interface. The * character indicates a continuation of addresses for the same session. For aggregated Ethernet interfaces, the output of the **summary (pic | port | slot)** options prefixes the interface name with **ae0:**.
Count	Count of subscribers displayed for each PIC, port, or slot when those options are specified with the **summary** option. For an aggregated Ethernet configuration, the total subscriber count does not equal the sum of the individual PIC, port, or slot counts, because each subscriber can be in more than one aggregated Ethernet link.
Total Subscribers	Total number of subscribers for all physical interfaces, all PICS, all ports, or all LS:RI slots.
IP Address/VLAN ID	Subscriber IP address or VLAN ID associated with the subscriber in the form *tpid.vlan-id*
User Name	Name of subscriber.
LS:RI	Logical system and routing instance associated with the subscriber.

Sample Output

show subscribers summary

```
user@host> show subscribers summary

Subscribers by State
 Init      3
 Configured    2
 Active   183
 Terminating    2
 Terminated    1

 TOTAL        191

Subscribers by Client Type
 DHCP     107
 PPP       76
 VLAN         8

 TOTAL        191
```

show subscribers summary all

```
user@host> show subscribers summary all
Subscribers by State
 Init         3
 Configured       2
 Active      183
 Terminating       2
 Terminated       1

 TOTAL          191

Subscribers by Client Type
 DHCP      107
 PPP          76
 VLAN          8

 TOTAL         191

Subscribers by LS:RI
 default:default   1
 default:ri1     28
 default:ri2     16
 ls1:default    22
 ls1:riA     38
 ls1:riB     44
 logsysX:routinstY  42

 TOTAL     191
```

show subscribers summary physical-interface

```
user@host> show subscribers summary physical-interface ge-1/0/0
Subscribers by State
   Active: 3998
   Total: 3998

Subscribers by Client Type
   DHCP: 3998
```

```
   Total: 3998

Subscribers by LS:RI
   default:default: 3998
   Total: 3998
```

show subscribers summary physical-interface pic

```
user@host> show subscribers summary physical-interface ge-0/2/0 pic
Subscribers by State
   Active: 4825
   Total: 4825

Subscribers by Client Type
   DHCP: 4825
   Total: 4825

Subscribers by LS:RI
   default:default: 4825
   Total: 4825
```

show subscribers summary physical-interface port

```
user@host> show subscribers summary physical-interface ge-0/3/0 port
Subscribers by State
   Active: 4825
   Total: 4825

Subscribers by Client Type
   DHCP: 4825
   Total: 4825

Subscribers by LS:RI
   default:default: 4825
   Total: 4825
```

show subscribers summary physical-interface slot

```
user@host> show subscribers summary physical-interface ge-2/0/0 slot
Subscribers by State
   Active: 4825
   Total: 4825

Subscribers by Client Type
   DHCP: 4825
   Total: 4825

Subscribers by LS:RI
   default:default: 4825
   Total: 4825
```

show subscribers summary pic

```
user@host> show subscribers summary pic
Interface           Count
ge-1/0              1000
ge-1/3              1000

Total Subscribers: 2000
```

show subscribers summary pic (Aggregated Ethernet Interfaces)

```
user@host> show subscribers summary pic
Interface           Count
ae0: ge-1/0         801
ae0: ge-1/3         801

Total Subscribers: 801
```

show subscribers summary port

```
user@host> show subscribers summary port
Interface           Count
ge-1                2000

Total Subscribers: 2000
```

show subscribers summary slot

```
user@host> show subscribers summary slot
Interface           Count
ge-1                2000

Total Subscribers: 2000
```

show subscribers summary terse

```
user@host> show subscribers summary terse
Interface           IP Address/VLAN ID   User Name         LS:RI
ge-1/3/0.1073741824   100                                   default:default
demux0.1073741824     100.0.0.10           WHOLESALER-CLIENT default:default
demux0.1073741825     101.0.0.3            RETAILER1-CLIENT  test1:retailer1
demux0.1073741826     102.0.0.3            RETAILER2-CLIENT  test1:retailer2
```

PART 4

Index

Index

Symbols

A

B

T

U

V