

Junos[®] OS

Broadband Subscriber Management Wholesale User Guide

Published
2020-03-19

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

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners.

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

Junos[®] OS Broadband Subscriber Management Wholesale User Guide
Copyright © 2020 Juniper Networks, Inc. All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

About the Documentation | xiv

Documentation and Release Notes | xiv

Using the Examples in This Manual | xiv

 Merging a Full Example | xv

 Merging a Snippet | xvi

Documentation Conventions | xvi

Documentation Feedback | xix

Requesting Technical Support | xix

 Self-Help Online Tools and Resources | xx

 Creating a Service Request with JTAC | xx

1

Configuring DHCP Layer 3 Wholesale Networks

Subscriber Management DHCP Layer 3 Wholesale Overview | 2

Layer 2 and Layer 3 Wholesale Overview | 2

Wholesale Network Configuration Options and Considerations | 3

DHCP Layer 3 Wholesale Configuration Interface Support | 4

Layer 3 Wholesale Configuration DHCP Support | 4

Subscriber to Logical System and Routing Instance Relationship | 5

RADIUS VSAs and Broadband Subscriber Management Wholesale Configuration Overview | 6

Configuring DHCPv4 Layer 3 Wholesale Networks | 8

Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Topology and Configuration Elements | 8

DHCPv4 Layer 3 Wholesale Network Topology Overview | 10

Configuring Loopback Interfaces for the DHCPv4 Layer 3 Wholesale Solution | 11

Configuring VLANs for the DHCPv4 Layer 3 Wholesale Network Solution | 12

 Configuring Static Customer VLANs for the DHCPv4 Layer 3 Wholesale Network Solution | 12

 Configuring Dynamic VLANs for the DHCPv4 Layer 3 Wholesale Network Solution | 13

Configuring Access Components for the DHCP Layer 3 Wholesale Network Solution | 16

 Configuring RADIUS Server Access | 17

 Configuring a DHCP Wholesaler Access Profile | 17

Configuring DHCP Retailer Access Profiles | 19

Configuring Dynamic Profiles for the DHCPv4 Layer 3 Wholesale Network Solution | 20

Configuring a Wholesale Dynamic Profile for use in the DHCPv4 Solution | 20

Configuring a Dynamic Profile for use by a Retailer in the DHCPv4 Solution | 22

Configuring Separate Routing Instances for DHCPv4 Service Retailers | 23

Configure Default Forwarding Options for the DHCPv4 Wholesale Network Solution | 27

Example: Wholesaler Dynamic Profile for a DHCPv4 Wholesale Network | 29

Example: Retailer Dynamic Profile for a DHCPv4 Wholesale Network | 30

Example: Default Forwarding Options Configuration for the DHCPv4 Wholesale Network | 31

Example: Retailer Routing Instances for a DHCPv4 Wholesale Network | 32

Configuring DHCPv6 Layer 3 Wholesale Networks | 36

Broadband Subscriber Management DHCPv6 Layer 3 Wholesale Topology and Configuration Elements | 36

DHCPv6 Layer 3 Wholesale Network Topology Overview | 38

Configuring Loopback Interfaces for the DHCPv6 Layer 3 Wholesale Solution | 39

Configuring VLANs for the DHCPv6 Layer 3 Wholesale Network Solution | 40

Configuring Static Customer VLANs for the DHCPv6 Layer 3 Wholesale Network Solution | 40

Configuring Dynamic Customer VLANs for the DHCPv6 Layer 3 Wholesale Network Solution | 41

Configuring Access Components for the DHCP Layer 3 Wholesale Network Solution | 44

Configuring RADIUS Server Access | 44

Configuring a DHCP Wholesaler Access Profile | 45

Configuring DHCP Retailer Access Profiles | 46

Configuring Dynamic Profiles for the DHCPv6 Layer 3 Wholesale Network Solution | 47

Configuring a Wholesale Dynamic Profile for use in the DHCPv6 Solution | 48

Configuring a Dynamic Profile for use by Each Retailer in the DHCPv6 Solution | 49

Configuring Separate Routing Instances for DHCPv6 Service Retailers | 51

Configuring Address Server Elements for the DHCPv6 Layer 3 Wholesale Solution | 52

Configuring a DHCPv6 Address Assignment Pool | 52

Configuring Extended DHCPv6 Local Server | 54

Example: Retailer Dynamic Profile for a DHCPv6 Wholesale Network | 56

Example: Retailer Routing Instances for a DHCPv6 Wholesale Network | 57

Example: DHCPv6 Address Assignment Pool That Provides Full 128-bit IPV6 Addresses for a DHCPv6 Wholesale Network | 58

Example: DHCPv6 Address Assignment Pool That Provides 74-bit IPV6 Prefixes for a DHCPv6 Wholesale Network | 59

Example: Extended DHCPv6 Local Server for a DHCPv6 Wholesale Network | 59

2

Configuring PPPoE Layer 3 Wholesale Networks

Subscriber Management PPPoE Wholesale Overview | 62

Layer 2 and Layer 3 Wholesale Overview | 62

PPPoE Layer 3 Wholesale Configuration Interface Support | 63

Subscriber to Logical System and Routing Instance Relationship | 64

RADIUS VSAs and Broadband Subscriber Management Wholesale Configuration Overview | 64

Configuring PPPoE Layer 3 Wholesale Networks | 66

Broadband Subscriber Management PPPoE Layer 3 Wholesale Topology and Configuration Elements | 66

PPPoE Layer 3 Wholesale Network Topology Overview | 68

Configuring Loopback Interfaces for the PPPoE Layer 3 Wholesale Solution | 69

Configuring Static Customer VLANs for the PPPoE Layer 3 Wholesale Network Solution | 70

Configuring Access Components for the PPPoE Wholesale Network Solution | 72

Configuring RADIUS Server Access | 72

Configuring a PPPoE Wholesaler Access Profile | 73

Configuring PPPoE Retailer Access Profiles | 74

Configuring Dynamic Profiles for the PPPoE Layer 3 Wholesale Network Solution | 75

Configuring a Wholesale Dynamic Profile for use in the PPPoE Solution | 76

Configuring Separate Routing Instances for PPPoE Service Retailers | 78

Example: Wholesaler Dynamic Profile for a PPPoE Wholesale Network | 79

Example: Retailer Routing Instances for a PPPoE Wholesale Network | 80

3

Configuring Layer 2 Wholesale Networks

Subscriber Management Layer 2 Wholesale Overview | 83

Layer 2 and Layer 3 Wholesale Overview | 83

Wholesale Network Configuration Options and Considerations | 84

RADIUS VSAs and Broadband Subscriber Management Wholesale Configuration Overview | 85

Extensible Subscriber Services Manager | 87

Extensible Subscriber Services Manager Overview | 87

Understanding the Dictionary File | 88

Configuring Layer 2 Wholesale Networks | 89

Broadband Subscriber Management Layer 2 Wholesale Topology and Configuration Elements | 89

Layer 2 Wholesale Network Topology Overview | 91

Configuring a Retail Dynamic Profile for Use in the Layer 2 Wholesale Solution | 92

Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution | 94

Configuring VLAN Interfaces for the Layer 2 Wholesale Solution | 97

Configuring Encapsulation for Layer 2 Wholesale VLAN Interfaces | 98

Configuring NNI ISP-Facing Interfaces for the Layer 2 Wholesale Solution | 100

Configuring Direct ISP-Facing Interfaces for the Layer 2 Wholesale Solution | 101

Configuring Separate Access Routing Instances for Layer 2 Wholesale Service Retailers | 102

Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers | 105

Configuring Access Components for the Layer 2 Wholesale Network Solution | 108

Configuring RADIUS Server Access | 108

Configuring a Layer 2 Wholesaler Access Profile | 109

Example: Retailer Dynamic Profile for a Layer 2 Wholesale Network | 110

Example: Access Interface for a Layer 2 Wholesale Network | 111

Example: Retailer Access Routing Instances for a Layer 2 Wholesale Network | 112

Example: Retailer NNI ISP-Facing Interfaces for a Layer 2 Wholesale Network | 113

Example: Retailer Direct ISP-Facing Interface for a Layer 2 Wholesale Network | 114

Configuring ANCP-Triggered Layer 2 Wholesale Services

ANCP-Triggered Layer 2 Wholesale Service Overview | 116

Layer 2 Wholesale with ANCP-Triggered VLANs Overview | 116

RADIUS Authorization for ANCP-Triggered VLANs | 119

Instantiation of an ANCP-Triggered, Autosensed, Dynamic VLAN | 119

Weighted Load Balancing for Subscriber Sessions over Eligible Core-Facing Physical Interfaces | 120

RADIUS Interim Accounting Updates | 121

Removal of the Layer 2 Wholesale Service | 122

Interactions Between In-Band and Out-of-Band VLAN Autosensing | 123

Migration of Subscriber Ownership from Wholesaler to Retailer	125
Migration of Subscriber Ownership from Retailer to Wholesaler	126
Migration of Subscriber Ownership Between Retailers	127
Modification of the Access Line Identifier or Port VLAN Identifier	127
Disconnecting PPPoE Sessions and Automatically Attempting Reconnection as Layer 2 Wholesale Sessions	128
Consequences of a State Transition in the Access-Facing Physical Interface	129
Consequences of a State Transition from Up to Down in the Core-Facing Physical Interface	131
Consequences of a State Transition from Down to Up in the Core-Facing Physical Interface	132
Loss of ANCP TCP Adjacency	133

Configuring ANCP-Triggered Layer 2 Wholesale Services | 136

Configuring ANCP Neighbors	136
Configuring the ANCP Agent for ANCP-Triggered, Autosensed Dynamic VLANs	138
Configuring a Username for Authentication of Out-of-Band Triggered Dynamic VLANs	139
Configuring Out-of-Band ANCP Messages to Trigger Dynamic VLAN Instantiation	141
Triggering ANCP OAM to Simulate ANCP Port Down and Port Up Messages	142
Configuring the ANCP Agent to Dampen the Effects of Short-Term Adjacency Losses	144
Reestablishing Pending Access Line Sessions for Layer 2 Wholesale	145
Configuring Multiple Non-Overlapping VLAN Ranges for Core-Facing Physical Interfaces	146
Clearing ANCP Access Loops	147

Configuring Flat-File Accounting for Layer 2 Wholesale Services | 149

Flat-File Accounting Overview	149
Configuring Flat-File Accounting for Layer 2 Wholesale	153
Configuring Flat-File Accounting for Extensible Subscriber Services Management	158
Configuring Service Accounting in Local Flat Files	163

Configuring Five-Level and Four-Level Heterogeneous Networks | 168

Five-Level and Four-Level Heterogeneous Networks	168
CoS Node Shaping in Four-Level and Five-Level Heterogeneous Networks	168
CuTTB Use Case Topology and CoS Hierarchy	173
FTTB/FTTH Use Case Topology and CoS Hierarchy	178
Automatic Creation of Business Subscriber Interface Sets	183
How to Configure the Automatic Creation of Business Subscriber Interface Sets	184

Dynamic Level 2 and Level 3 Interface Set Naming with Predefined Variables | 185

Predefined Variable Default for the Level 2 Node Interface Set | 187

Predefined Variable Default for the Level 3 Node Interface Set | 187

OLT Migration to Using PON TLVs Instead of DSL TLVs | 190

Support for OLT Migration to PON TLVs | 190

How to Configure Preference for DSL or PON TLVs When an OLT Sends Both | 191

Configuration Statements and Operational Commands

Configuration Statements | 194

accept | 198

accept-out-of-band | 200

access-profile | 201

access-profile (Dynamic VLAN) | 203

access-profile (Dynamic Stacked VLAN) | 204

active-server-group | 205

address | 207

address-assignment (Address-Assignment Pools) | 211

adjacency-loss-hold-time (ANCP) | 214

ancp | 215

authentication | 217

authentication (DHCP Local Server) | 219

authentication (DHCP Relay Agent) | 221

authentication-order | 223

auto-configure | 225

auto-configure-trigger interface (ANCP) | 227

backup-on-failure (Accounting Options) | 228

circuit-id (VLAN Authentication Username) | 229

cleanup-interval (Accounting Options) | 230

compress (Accounting Options) | 231

connectivity-type | 232

core-facing | 233

demux0 (Dynamic Interface) | 234

demux-options (Dynamic Interface) | 236

demux-source (Dynamic IP Demux Interface) | 237

demux-source (Dynamic Underlying Interface) | 238

demux-source (Underlying Interface) | 239

dhcp-attributes (Address-Assignment Pools) | 240

dhcp-local-server | 247

dhcp-relay | 264

dhcpv6 (DHCP Local Server) | 283

dynamic-profile (DHCP Local Server) | 292

dynamic-profile (DHCP Relay Agent) | 294

dynamic-profile (Dynamic PPPoE) | 296

dynamic-profile (Stacked VLAN) | 298

dynamic-profile (VLAN) | 299

dynamic-profiles | 300

egress-stats (Flat-File Accounting Options) | 314

encapsulation (Dynamic Interfaces) | 316

exclude (RADIUS Attributes) | 319

family (Address-Assignment Pools) | 327

family (Dynamic Demux Interface) | 329

family (Dynamic PPPoE) | 331

family (Dynamic Standard Interface) | 333

fields (Flat-File Accounting Options) | 336

file (Flat-File Accounting Options) | 339

flat-file-profile (Accounting Options) | 340

flat-file-profile (Extensible Subscriber Services) | 342

flexible-vlan-tagging | 343

format (Flat-File Accounting Options) | 344

forwarding-options | 345

general-param (Flat-File Accounting Options) | 348

group (DHCP Local Server) | 350

group (DHCP Relay Agent) | 356

ingress-stats (Flat-File Accounting Options) | 362

inner-vlan-id (Dynamic VLANs) | 364

inner-vlan-id-swap-ranges | 365

input-vlan-map (Dynamic Interfaces) | 366

instance-role | 367

instance-type | 369

interface (DHCP Local Server) | 372

interface (DHCP Relay Agent) | 375

interface (Dynamic Routing Instances) | 378

interface (Routing Instances) | 379

interface-mac-limit (VPLS) | 380

interfaces (Static and Dynamic Subscribers) | 382

interval (Flat-File Accounting Options) | 389

ip-address-first | 390

keepalives (Dynamic Profiles) | 391

l2-stats (Flat-File Accounting Options) | 393

mac-validate (Dynamic IP Demux Interface) | 394

multicast-replication | 395

neighbor (Define ANCP) | 397

no-local-switching | 398

no-tunnel-services | 399

overall-packet (Flat-File Accounting Options) | 401

output-vlan-map (Dynamic Interfaces) | 403

pap (Dynamic PPP) | 404

pool (Address-Assignment Pools) | 405

pool-match-order | 407

pop (Dynamic VLANs) | 408

pppoe-options (Dynamic PPPoE) | 409

pppoe-underlying-options (Static and Dynamic Subscribers) | 410

ppp-options (Dynamic PPP) | 411

prefix (Address-Assignment Pools) | 413

profile (Access) | 414

protocols | 421

proxy-arp | 425

proxy-arp (Dynamic Profiles) | 426

push (Dynamic VLANs) | 427

push-backup-to-master (Accounting Options) | 428

radius (Access Profile) | 429

radius-server | 433

range (Address-Assignment Pools) | 439

ranges (Dynamic VLAN) | 440

remote-id (VLAN Authentication Username) | 441

route-distinguisher | 442

routing-instances (Dynamic Profiles) | 446

routing-instances (Multiple Routing Entities) | 448

schema-version (Flat-File Accounting Options) | 450

secret (RADIUS) | 451

server (Dynamic PPPoE) | 452

server-group | 453

site (VPLS Multihoming for FEC 128) | 455

site-identifier (VPLS) | 456

site-range | 457

stacked-vlan-ranges | 458

stacked-vlan-tagging | 459

traceoptions (DHCP) | 460

underlying-interface (demux0) | 463

underlying-interface (Dynamic PPPoE) | 465

unit | 466

unit (Dynamic Demux Interface) | 477

unit (Dynamic Profiles Standard Interface) | 479

unnumbered-address (Dynamic PPPoE) | 483

unnumbered-address (Dynamic Profiles) | 484

unnumbered-address (Ethernet) | 486

username-include (Interfaces) | 487

user-prefix (DHCP Local Server) | 489

vlan-id (Dynamic VLANs) | 491

vlan-model | 492

vlan-ranges | 493

vlan-tags | 495

vlan-tags (Stacked VLAN Tags) | 497

vpls (Routing Instance) | 499

vrf-export | 502

vrf-import | 504

vrf-target | 506

Operational Commands | 508

clear ancp access-loop | 510

clear ancp neighbor | 512

clear dhcp relay binding | 514

clear dhcp relay statistics | 518

clear dhcp server binding | 521

clear dhcp server statistics | 525

clear dhcpv6 server binding | 527

clear dhcpv6 server statistics | 530

clear network-access aaa subscriber | 532

request ancp oam port-down | 534

request ancp oam port-up | 536

request auto-configuration reconnect-pending | 538

show ancp neighbor | 539

show ancp subscriber | 549

show auto-configuration out-of-band | 559

show dhcp relay binding | 565

show dhcp relay statistics | 572

show dhcp server binding | 577

show dhcp server statistics | 587

show dhcpv6 server binding | 592

show dhcpv6 server statistics | 600

show extensible-subscriber-services counters | 604

show extensible-subscriber-services debug-information | 606

show extensible-subscriber-services dictionary | 607

show extensible-subscriber-services dictionary attributes | 612

show extensible-subscriber-services dictionary services | 615

show extensible-subscriber-services sessions | 619

show extensible-subscriber-services service | 621

show interfaces (Fast Ethernet) | 623

show interfaces | 646

show interfaces (Loopback) | 753

show interfaces (PPPoE) | **763**

show interfaces demux0 (Demux Interfaces) | **776**

show interfaces filters | **788**

show interfaces l2-routing-instance | **791**

show interfaces routing | **793**

show interfaces routing-instance | **801**

show network-access aaa statistics | **804**

show network-access aaa statistics authentication | **817**

show network-access aaa subscribers | **821**

show network-access address-assignment pool | **827**

show ppp interface | **829**

show subscribers | **849**

show subscribers summary | **896**

show vpls connections | **905**

show vpls flood event-queue | **922**

show vpls flood instance | **924**

show vpls flood route | **927**

show vpls mac-table | **930**

show vpls statistics | **937**

About the Documentation

IN THIS SECTION

- Documentation and Release Notes | xiv
- Using the Examples in This Manual | xiv
- Documentation Conventions | xvi
- Documentation Feedback | xix
- Requesting Technical Support | xix

Use this guide to understand how wholesaling enables service providers to resell broadband services and enable other providers to deploy their own services over the incumbent network. This guide discusses Layer 3 wholesale networks (DHCP and PPPoE) and Layer 2 wholesale networks (including ANCP-triggered wholesale services).

Documentation and Release Notes

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

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

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

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.xml; }
```

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

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

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

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

For more information about the **load** command, see [CLI Explorer](#).

Documentation Conventions

[Table 1 on page xvii](#) defines notice icons used in this guide.

Table 1: Notice Icons







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

Table 2 on page xvii 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
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies guide names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS CLI User Guide</i> RFC 1997, <i>BGP Communities Attribute</i>

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

Convention	Description	Examples
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none">• To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level.• The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	stub <default-metric <i>metric</i>>;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (<i>string1</i> <i>string2</i> <i>string3</i>)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	community name members [<i>community-ids</i>]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		


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

Convention	Description	Examples
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> In the Logical Interfaces box, select All Interfaces. To cancel the configuration, click Cancel.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

We encourage you to provide feedback so that we can improve our documentation. You can use either of the following methods:

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the [Juniper Networks TechLibrary](#) site, and do one of the following:



- Click the thumbs-up icon if the information on the page was helpful to you.
- Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are

covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

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

Self-Help Online Tools and Resources

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

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

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

Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit <https://myjuniper.juniper.net>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see <https://support.juniper.net/support/requesting-support/>.

1

PART

Configuring DHCP Layer 3 Wholesale Networks

[Subscriber Management DHCP Layer 3 Wholesale Overview](#) | 2

[Configuring DHCPv4 Layer 3 Wholesale Networks](#) | 8

[Configuring DHCPv6 Layer 3 Wholesale Networks](#) | 36

Subscriber Management DHCP Layer 3 Wholesale Overview

IN THIS CHAPTER

- [Layer 2 and Layer 3 Wholesale Overview | 2](#)
- [Wholesale Network Configuration Options and Considerations | 3](#)
- [DHCP Layer 3 Wholesale Configuration Interface Support | 4](#)
- [Layer 3 Wholesale Configuration DHCP Support | 4](#)
- [Subscriber to Logical System and Routing Instance Relationship | 5](#)
- [RADIUS VSAs and Broadband Subscriber Management Wholesale Configuration Overview | 6](#)

Layer 2 and Layer 3 Wholesale Overview

In general, wholesaling broadband services allows service providers to resell broadband services and allows other providers to deploy their own services over the incumbent network. There are different methods to partitioning an access network for resale. The two most common approaches are based on either Layer 2 or Layer 3 information. Wholesale access is the process by which the access network provider (the *wholesaler*) partitions the access network into separately manageable and accountable subscriber segments for resale to other network providers (or *retailers*).

In a Layer 3 wholesale configuration, you partition the wholesaler access network at the network layer or the subscriber IP component by associating the IP component with a distinct Layer 3 domain. In a Layer 2 wholesale configuration, you partition the access network at the subscriber circuit or customer VLAN (C-VLAN) by backhauling the connection through the service provider backbone network to the subscribing retailer network where the access traffic can be managed at higher layers.

In a Junos OS Dynamic Host Configuration Protocol (DHCP) or Point-to-Point Protocol over Ethernet (PPPoE) subscriber access configuration, wholesale partitioning is accomplished through the use of logical systems and routing instances within the router. Logical systems offer a stricter partitioning of routing resources than routing instances. The purpose behind the use of logical systems is to distinctly partition the physical router into separate administrative domains. This partitioning enables multiple providers to administer the router simultaneously, with each provider having access only to the portions of the configuration relevant to their logical system. Junos OS supports up to 15 named logical systems in addition

to the default logical system (that is, **inet.0**). Unless otherwise specified in configuration, all interfaces belong to the default logical system.

NOTE: This Junos OS release supports the use of only the default logical system. Partitioning currently occurs through the use of separate routing instances.

A logical system can have one or more routing instances. Typically used in Layer 3 VPN scenarios, a routing instance does not have the same level of administrative separation as a logical system because it does not offer administrative isolation. However, the routing instance defines a distinct routing table, set of routing policies, and set of interfaces.

RELATED DOCUMENTATION

[Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Topology and Configuration Elements | 8](#)

[Broadband Subscriber Management PPPoE Layer 3 Wholesale Topology and Configuration Elements | 66](#)

[Broadband Subscriber Management Layer 2 Wholesale Topology and Configuration Elements | 89](#)

Wholesale Network Configuration Options and Considerations

You can configure a wholesale network any number of ways using Juniper Networks hardware and Junos OS software. The general configuration options, and considerations for each, are provided in the following table:

Wholesale Configuration Options	Considerations
Fully Static (all interfaces, VLANs, and routing instances are configured statically)	Providing more control over retailer space and access, this option is more labor intensive and can require more detailed planning of the network, address allocation, and so on.
Static VLANs and Dynamic Demux Interfaces	Service VLANs are created statically and must be managed. Demux interfaces are dynamically created over the service VLANs. This option uses more logical interfaces; one for each VLAN and one for each dynamic demux interface that runs over each VLAN.

Wholesale Configuration Options	Considerations
Dynamic VLANs Only (dedicated customer VLANs for each subscriber)	<p>Dynamic (auto-sensed) VLANs are authenticated and installed in the correct non-default routing instance before DHCP is instantiated. This method helps to conserve logical interfaces by avoiding the need for additional logical interfaces being created for each demux interface.</p> <p>NOTE: In a customer VLAN model, each VLAN functions on a 1:1 basis for each customer (in this case, per household).</p>
Dynamic VLANs and Dynamic Demux Interfaces	<p>Allows for the greatest ease of use and flexibility in configuring subscribers, by enabling access over a service VLAN and targetting more service levels over individual, dynamically-created demux interfaces over the service VLAN. This option uses more logical interfaces; one for each VLAN and one for each demux interface that runs over each VLAN.</p>

DHCP Layer 3 Wholesale Configuration Interface Support

DHCP Layer 3 wholesale currently supports only the use of IP demux interfaces.

For general additional information about configuring IP demux interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*.

RELATED DOCUMENTATION

Junos OS Network Interfaces Library for Routing Devices

Subscriber Interfaces and Demultiplexing Overview .

Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles.

Configuring a Subscriber Interface Using a Set of Static IP Demux Interfaces.

Layer 3 Wholesale Configuration DHCP Support

DHCP Layer 3 wholesale supports the following DHCP configuration options:

- DHCP Relay
- DHCP Relay Proxy
- DHCP Local Server

NOTE: All routing instances within the same wholesale network must use the same DHCP configuration option.

For additional information about any of these DHCP options, see the *AAA Service Framework Overview*.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview.

DHCP Relay Proxy Overview.

Understanding Differences Between Legacy DHCP and Extended DHCP.

Subscriber to Logical System and Routing Instance Relationship

As subscriber sessions are established, subscriber to logical system/routing instance memberships are established by the AAA framework configured for the default logical system. When configuring Layer 3 wholesaling, you typically configure global (wholesale) information within the default (master) logical system and default routing instance. Incoming subscribers must then be authenticated, but this authentication can be handled in one of two ways:

- Single (wholesaler only) authentication—Incoming subscribers are authenticated by the wholesaler RADIUS server. After authentication, the subscribers are assigned values specified by dynamic profiles (routing instances, interfaces, and any configuration values) specific to a particular retailer.
- Dual (wholesaler and retailer) authentication—Sometimes referred to as *double-dip authentication*. Incoming subscribers are initially authenticated by RADIUS using the wholesale configuration. Authenticated subscribers are then redirected to other routing instances associated with individual retailer network space. When you redirect subscribers, and those subscribers are to be authenticated by AAA servers owned by individual retailers, the subscribers must be authenticated again by the AAA servers before they are provided an address and any dynamic profile values are assigned. After reauthentication, however, the subscribers are managed normally using any values specific to the retailer routing instance to which they are assigned.

RELATED DOCUMENTATION

See *Routing Instances Overview*.

RADIUS VSAs and Broadband Subscriber Management Wholesale Configuration Overview

You can use RADIUS to assign various values through the use of dynamic variables within dynamic profiles. However, the configuration of at least one of the two VSAs described in [Table 3 on page 6](#) is required for a wholesale network to function.

Table 3: Required Juniper Networks VSAs for the Broadband Subscriber Management Wholesale Network Solution

Attribute Number	Attribute Name	Description	Value
26-1	LSRI-Name	Client logical system/routing instance membership name. Allowed only from RADIUS server for “default” logical system/routing instance membership.	string: logical system:routing instance
26-25	Redirect-LSRI-Name	Client logical system/routing instance membership name indicating to which logical system/routing instance membership the request is redirected for user authentication.	string: logical system:routing instance

Specifying the **\$junos-routing-instance** dynamic variable in a dynamic profile triggers a RADIUS access-accept response of either the LSRI-Name VSA or the Redirect-LSRI-Name VSA. Returning an LSRI-Name attribute in the access-accept response provides the logical system and routing instance in which the logical interface is to be created and the router updates the session database with the specified routing instance value. Returning a Redirect-LSRI-Name attribute in the access-accept response results in the router immediately sending a second access-request message (sometimes referred to as a *double-dip*) to the RADIUS server specified by the logical system:routing instance attribute specified by the Redirect-LSRI-Name VSA.

NOTE: Attributes returned as a result of a second access-request message to the logical system/routing instance membership specified by the Redirect-LSRI-Name VSA override any prior attributes returned by initial access-accept responses to the default logical system/routing instance membership.

RELATED DOCUMENTATION

Juniper Networks VSAs Supported by the AAA Service Framework.

Configuring DHCPv4 Layer 3 Wholesale Networks

IN THIS CHAPTER

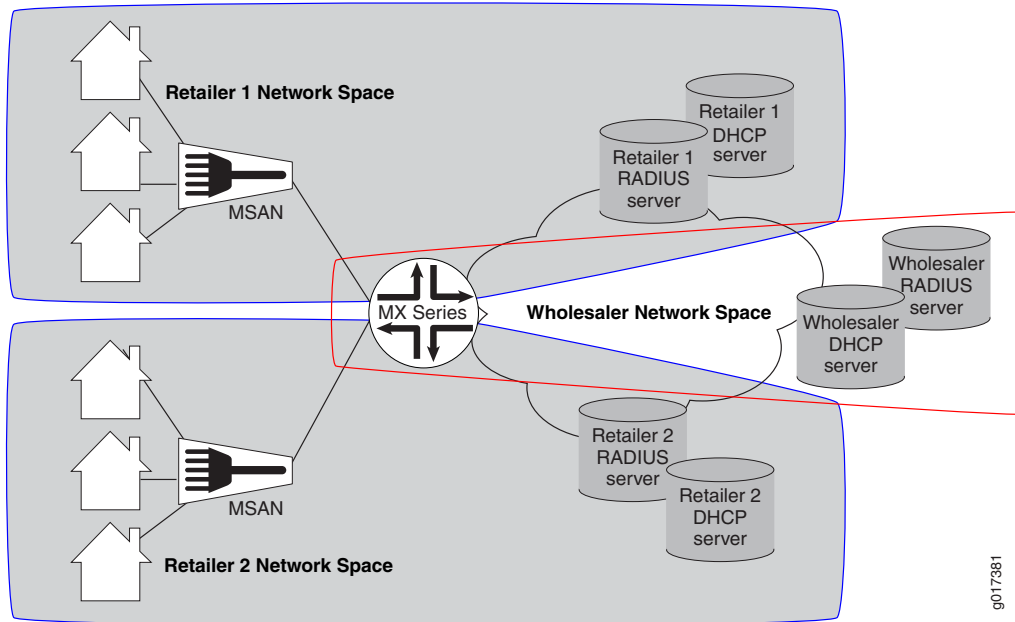
- Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Topology and Configuration Elements | 8
- DHCPv4 Layer 3 Wholesale Network Topology Overview | 10
- Configuring Loopback Interfaces for the DHCPv4 Layer 3 Wholesale Solution | 11
- Configuring VLANs for the DHCPv4 Layer 3 Wholesale Network Solution | 12
- Configuring Access Components for the DHCP Layer 3 Wholesale Network Solution | 16
- Configuring Dynamic Profiles for the DHCPv4 Layer 3 Wholesale Network Solution | 20
- Configuring Separate Routing Instances for DHCPv4 Service Retailers | 23
- Configure Default Forwarding Options for the DHCPv4 Wholesale Network Solution | 27
- Example: Wholesaler Dynamic Profile for a DHCPv4 Wholesale Network | 29
- Example: Retailer Dynamic Profile for a DHCPv4 Wholesale Network | 30
- Example: Default Forwarding Options Configuration for the DHCPv4 Wholesale Network | 31
- Example: Retailer Routing Instances for a DHCPv4 Wholesale Network | 32

Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Topology and Configuration Elements

The network topology for the subscriber management DHCPv4 Layer 3 wholesale solution includes configuring separate routing instances for individual retailers that use a portion of the router. This solution uses a DHCPv4 relay configuration. However, you can also implement DHCPv4 Relay Proxy or DHCPv4 Local Server configuration.

To explain the concept, but to limit complexity, this solution provides a configuration with one wholesaler and only two retailers. [Figure 1 on page 9](#) illustrates a basic Layer 3 wholesale topology model from which you can expand.

Figure 1: Basic Subscriber Management Layer 3 Wholesale Solution Topology



A DHCP Layer 3 wholesale network solution can use various combinations of the following configuration elements:

- Subscriber network VLAN configuration
- DHCPv4 configuration (DHCPv4 Relay, DHCPv4 Relay Proxy, or DHCPv4 Local Server)
- Addressing server or addressing server access configuration (if not using DHCPv4 Local Server)
- RADIUS server access configuration
- Dynamic profile configuration for default (wholesaler) access
- Dynamic profile configuration for retailer access (following subscriber redirection, if applicable)
- Routing instance configuration for individual retailers
- Group configuration and forwarding options for the network
- Core network configuration

RELATED DOCUMENTATION

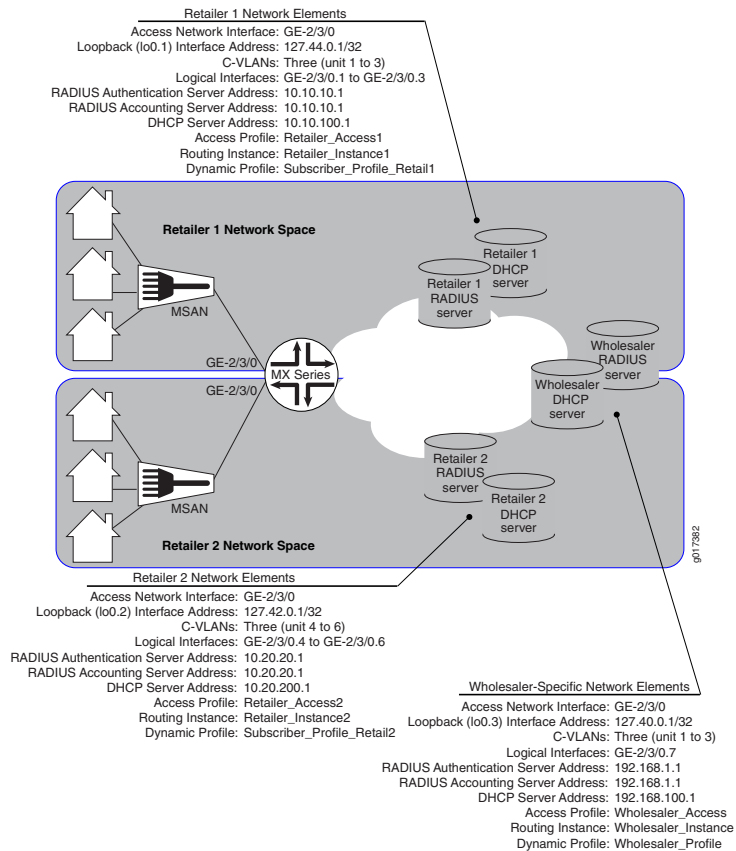
[Layer 2 and Layer 3 Wholesale Overview | 2](#)

[DHCPv4 Layer 3 Wholesale Network Topology Overview | 10](#)

DHCPv4 Layer 3 Wholesale Network Topology Overview

This configuration explains how to configure a simple DHCPv4 Layer 3 wholesale subscriber access network. This solution incorporates two retailers sharing resources on a wholesaler router. [Figure 2 on page 10](#) provides the reference topology for this configuration example.

Figure 2: DHCPv4 Layer 3 Wholesale Network Reference Topology



RELATED DOCUMENTATION

[Layer 2 and Layer 3 Wholesale Overview | 2](#)

[Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Topology and Configuration Elements | 8](#)

Configuring Loopback Interfaces for the DHCPv4 Layer 3 Wholesale Solution

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

To configure loopback interfaces:

1. Edit the loopback interface.

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

2. Edit the unit for the wholesale loopback interface.

```
[edit interfaces lo0]  
user@host# edit unit 3
```

3. Edit the loopback interface family that belongs to the wholesaler.

```
[edit interfaces lo0 unit 3]  
user@host# edit family inet
```

4. Specify the loopback interface address that belongs to the wholesaler.

```
[edit interfaces lo0 unit 3]  
user@host# set address 127.40.0.1/32
```

5. Edit the unit for a retail loopback interface to be assigned to the retailer.

```
[edit interfaces lo0]  
user@host# edit unit 1
```

6. Edit the loopback interface family that will be assigned to the retailer.

```
[edit interfaces lo0 unit 1]  
user@host# edit family inet
```

- Specify the loopback interface address that will be assigned to the retailer.

```
[edit interfaces lo0 unit 1]
user@host# set address 127.42.0.1/32
```

- Repeat steps 5 through 7 for additional retailers, making sure to use unique unit and address values for each retailer loopback interface.

RELATED DOCUMENTATION

Junos OS Network Interfaces Library for Routing Devices

Configuring VLANs for the DHCPv4 Layer 3 Wholesale Network Solution

IN THIS SECTION

- [Configuring Static Customer VLANs for the DHCPv4 Layer 3 Wholesale Network Solution | 12](#)
- [Configuring Dynamic VLANs for the DHCPv4 Layer 3 Wholesale Network Solution | 13](#)

You can configure either static or dynamic customer VLANs for use in the DHCPv4 wholesale network solution.

Configuring Static Customer VLANs for the DHCPv4 Layer 3 Wholesale Network Solution

In this example configuration, the access interface (**ge-2/3/0**) connects to a device (that is, a DSLAM) on the access side of the network. You can define static VLANs for use by the access network subscribers.

To configure the static VLANs:

- Edit the access side interface.

```
[edit]
user@host# edit interfaces ge-2/3/0
```

- Specify the use of stacked VLAN tagging.


```
[edit interfaces ge-2/3/0]
user@host# set stacked-vlan-tagging
```

3. Edit the interface unit for the first VLAN.

```
[edit interfaces ge-2/3/0]
user@host# edit unit 1
```

4. Define the VLAN tags for the first VLAN.

```
[edit interfaces ge-2/3/0 unit 1]
user@host# set vlan-tags outer 3 inner 1
```

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

```
[edit interfaces ge-2/3/0 unit 1]
user@host# set demux-source inet
```

6. Edit the family for the first VLAN.

```
[edit interfaces ge-2/3/0 unit 1]
user@host# edit family inet
```

7. (Optional) Define the unnumbered address and the preferred source address for the first VLAN.

```
[edit interfaces ge-2/3/0 unit 1 family inet]
user@host# set unnumbered-address lo0.1 preferred-source-address 127.44.0.1
```

8. Repeat steps 2 through 7 for additional VLAN interface units.

Configuring Dynamic VLANs for the DHCPv4 Layer 3 Wholesale Network Solution

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# edit 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 **[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 **[interfaces]** hierarchy level.

- h. Access the family type.

```
[edit dynamic-profiles VLAN-PROF interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit"]
user@host# edit 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. (Optional) 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 127.33.0.1
```

2. Associate the dynamic profile with the interface on which the dynamic VLANs will be created.

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

```
[edit interfaces]
user@host# edit interfaces ge-2/3/0
```

- b. Specify the use of stacked VLAN tagging.

```
[edit interfaces ge-2/3/0]
user@host# set stacked-vlan-tagging
```

- c. Specify that you want to automatically configure VLAN interfaces.

```
[edit interfaces ge-2/3/0]
user@host# edit auto-configure
```

- d. Specify that you want to configure stacked VLANs.

```
[edit interfaces ge-2/3/0 auto-configure]
user@host# edit stacked-vlan-ranges
```

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

```
[edit interfaces ge-2/3/0 auto-configure stacked-vlan-ranges]
user@host# set dynamic-profile VLAN-PROF
```

- f. Repeat steps a through e for any other interfaces that you want to use for creating VLANs.

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

```
[edit interfaces ge-2/3/0 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. 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 **1-3** (enabling a range from 1 through 3 for the inner stacked VLAN ID).

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

Configuring Access Components for the DHCP Layer 3 Wholesale Network Solution

IN THIS SECTION

- [Configuring RADIUS Server Access | 17](#)
- [Configuring a DHCP Wholesaler Access Profile | 17](#)
- [Configuring DHCP Retailer Access Profiles | 19](#)

When configuring a wholesale network, you must configure several components globally. This configuration provides access to RADIUS servers that you want the wholesaler and any configured retailers to use globally. The access configuration includes the following general steps:

Configuring RADIUS Server Access

You can globally define any RADIUS servers in your network that either the wholesale access profile or retailer access profile can use. After you define the global RADIUS servers, you can specify specific RADIUS servers within individual access profiles.

To define RADIUS servers for profile access:

1. Access the **[edit access radius-server]** hierarchy level.

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

2. Specify the address and secret for any RADIUS servers in the network.

```
[edit access radius-server]
user@host# set 192.168.10.1 secret $ABC123$ABC123$ABC123
user@host# set 10.10.10.1 secret $ABC123$ABC123
```

SEE ALSO

| *RADIUS Servers and Parameters for Subscriber Access*

Configuring a DHCP Wholesaler Access Profile

You must define the network and interface over which you want subscribers to initially access the network with a wholesale access profile. When a subscriber attempts to access the network, the access profile provides initial access information including authentication and accounting values that the router uses for the accessing subscriber.

To define a wholesale access profile:

1. Create the wholesale access profile.

```
[edit]
user@host# edit access-profile Wholesaler_Access
```

- Specify the authentication methods for the profile and the order in which they are used.

```
[edit access profile Wholesaler1]
user@host# set authentication-order radius password
```

- Specify that you want to configure RADIUS support.

```
[edit access profile Wholesaler1]
user@host# edit radius
```

- Specify the IP address of the RADIUS server used for authentication.

```
[edit access profile Wholesaler1 radius]
user@host# set authentication-server 192.168.10.1
```

- Specify the IP address of the RADIUS server used for accounting.

```
[edit access profile Wholesaler1 radius]
user@host# set accounting-server 192.168.10.1
```

- Configure any desired options for the RADIUS server.

See *RADIUS Servers and Parameters for Subscriber Access*.

- Configure subscriber accounting (RADIUS accounting).

See *Configuring Per-Subscriber Session Accounting*.

SEE ALSO

Configuring Authentication and Accounting Parameters for Subscriber Access

Specifying the Authentication and Accounting Methods for Subscriber Access

RADIUS Servers and Parameters for Subscriber Access

Configuring Per-Subscriber Session Accounting

Configuring DHCP Retailer Access Profiles

In this solution, subscribers are redirected to a networking space used by a specific retailer and defined by a unique routing instance. This method requires that you define the network and interface over which you want subscribers to access the network after being redirected by the wholesale access profile.

To define a retailer access profile:

1. Create the retailer access profile.

```
[edit]
user@host# edit access-profile Retailer_Access1
```

2. Specify the authentication methods for the profile and the order in which they are used.

```
[edit access profile Retailer1]
user@host# set authentication-order radius password
```

3. Specify that you want to configure RADIUS support.

```
[edit access profile Retailer1]
user@host# edit radius
```

4. Specify the IP address of the RADIUS server used for authentication.

```
[edit access profile Retailer1 radius]
user@host# set authentication-server 10.10.10.1
```

5. Specify the IP address of the RADIUS server used for accounting.

```
[edit access profile Retailer1 radius]
user@host# set accounting-server 10.10.10.1
```

6. Configure any desired options for the RADIUS server.

See *RADIUS Servers and Parameters for Subscriber Access*.

7. Configure subscriber accounting (RADIUS accounting).

See *Configuring Per-Subscriber Session Accounting*.

SEE ALSO

[*Configuring Authentication and Accounting Parameters for Subscriber Access*](#)[*Specifying the Authentication and Accounting Methods for Subscriber Access*](#)[*RADIUS Servers and Parameters for Subscriber Access*](#)[*Configuring Per-Subscriber Session Accounting*](#)

Configuring Dynamic Profiles for the DHCPv4 Layer 3 Wholesale Network Solution

IN THIS SECTION

- [Configuring a Wholesale Dynamic Profile for use in the DHCPv4 Solution | 20](#)
- [Configuring a Dynamic Profile for use by a Retailer in the DHCPv4 Solution | 22](#)

A dynamic profile is a set of characteristics, defined in a type of template, that you can use to provide services for broadband applications. These services are assigned dynamically to interfaces as they access the network. When configuring dynamic profiles for the DHCPv4 Layer 3 wholesale network, you can choose to configure one dynamic profile to address all incoming subscribers or you can configure individual dynamic profiles for use by the different network management groups (that is, the wholesaler and any retailers). In fact, you can create multiple dynamic profiles that you can use to roll out different services and selectively apply those dynamic profiles to different subscriber groups as necessary.

In this solution example, one dynamic profile is created for use by the wholesaler when subscribers initially access the network. Other dynamic profiles are created for the subscribers for each individual retailer to use after they are redirected to that retailer network space.

Configuring a Wholesale Dynamic Profile for use in the DHCPv4 Solution

You can configure a basic access profile to initially manage subscribers that access the network.

To configure a dynamic profile for use by the wholesaler:

1. Create a wholesale dynamic profile.

```
[edit]  
user@host# edit dynamic-profiles Wholesaler_Profile
```


2. Specify that you want to configure the **demux0** interface in the dynamic profile.

```
[edit dynamic-profiles Subscriber_Profile_Retail1]
user@host# edit interfaces demux0
```

3. Configure the unit for the **demux0** interface.

- a. Configure the variable for the unit number of the **demux0** interface.

The variable is dynamically replaced with the unit number that DHCP supplies when the subscriber logs in.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 demux0]
user@host# edit unit $junos-interface-unit
```

- b. Configure the variable for the underlying interface of the demux interfaces and specify the **\$junos-underlying-interface** variable.

The variable is dynamically replaced with the underlying interface that DHCP supplies when the subscriber logs in.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces demux0 unit "$junos-interface-unit"]
user@host# set demux-options underlying-interface $junos-underlying-interface
```

4. Configure the family for the demux interfaces.

- a. Specify that you want to configure the family.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces demux0 unit "$junos-interface-unit"]
user@host# edit family inet
```

- b. Configure the unnumbered address for the family.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 demux0 unit "$junos-interface-unit" family inet6]
user@host# set unnumbered-address lo0.0
```

- c. Configure the variable for the IPv4 address of the demux interface.

The variable is dynamically replaced with the IPv4 address that DHCP supplies when the subscriber logs in.

```
[edit dynamic-profiles business-profile interfaces demux0 unit "$junos-interface-unit"]
user@host# set demux-source $junos-subscriber-ip-address
```

Configuring a Dynamic Profile for use by a Retailer in the DHCPv4 Solution

To configure a dynamic profile for use with retailer access:

1. Create a retail dynamic profile.

```
[edit]
user@host# edit dynamic-profiles Subscriber_Profile_Retail1
```

2. Define the dynamic routing instance variable in the dynamic profile.

```
[edit dynamic-profiles Subscriber_Profile_Retail1]
user@host# edit routing-instances $junos-routing-instance
```

3. Set the dynamic interface variable for the dynamic routing instance.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 routing-instances "$junos-routing-instance"]
user@host# set interface $junos-interface-name
```

4. Specify that you want to configure the **demux0** interface in the dynamic profile.

```
[edit dynamic-profiles Subscriber_Profile_Retail1]
user@host# edit interfaces demux0
```

5. Configure the unit for the **demux0** interface.

- a. Configure the variable for the unit number of the **demux0** interface.

The variable is dynamically replaced with the unit number that DHCP supplies when the subscriber logs in.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 demux0]
user@host# edit unit $junos-interface-unit
```

- b. Configure the variable for the underlying interface of the demux interfaces and specify the **\$junos-underlying-interface** variable.

The variable is dynamically replaced with the underlying interface that DHCP supplies when the subscriber logs in.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces demux0 unit "$junos-interface-unit"]
user@host# set demux-options underlying-interface $junos-underlying-interface
```

6. Configure the family for the demux interfaces.

- a. Specify that you want to configure the family.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces demux0 unit "$junos-interface-unit"]
user@host# edit family inet
```

- b. Configure the unnumbered address for the family.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 demux0 unit "$junos-interface-unit" family inet6]
user@host# set unnumbered-address lo0.0
```

- c. Configure the variable for the IPv6 address of the demux interface.

The variable is dynamically replaced with the IPv6 address that DHCP supplies when the subscriber logs in.

```
[edit dynamic-profiles business-profile interfaces demu0 unit "$junos-interface-unit"]
user@host# set demux-source $junos-subscriber-ip-address
```

Configuring Separate Routing Instances for DHCPv4 Service Retailers

As the owner of the system, the wholesaler typically uses the default routing instance. You must create separate routing instances for each individual retailer to keep routing information for individual retailers separate and to define any servers and forwarding options specific to each retailer.

To define a retailer routing instance:

1. Create the retailer routing instance.

```
[edit]
user@host# edit routing-instances RetailerInstance1
```

2. Specify the routing instance type for the retailer.

```
[edit routing-instances "RetailerInstance1"]
user@host# set instance-type vrf
```

3. Specify the access profile that you want the routing instance to use.

```
[edit routing-instances "RetailerInstance1"]
user@host# set access-profile Retailer1
```

4. Specify the interface that faces the Retailer1 RADIUS server.

```
[edit routing-instances "RetailerInstance1"]
user@host# set interface ge-11/1/9.10
```

5. Specify the interface that faces the Retailer1 DHCP server.

```
[edit routing-instances "RetailerInstance1"]
user@host# set interface ge-11/1/10.100
```

6. Specify the loopback interface unit for this routing instance.

```
[edit routing-instances "RetailerInstance1"]
user@host# set interface lo0.1
```

NOTE: Loopback interfaces must be unique for each routing instance.

7. Access the DHCP Relay forwarding options hierarchy for the routing instance.

```
[edit routing-instances "RetailerInstance1"]
user@host# edit forwarding-options dhcp-relay
```

NOTE: The configuration for this wholesale solution uses DHCP Relay. However, you can also configure DHCP Proxy Relay or DHCP Local Server for the DHCP Layer 3 wholesale network.

8. Specify that you want to configure authentication options and use external AAA authentication services.

```
[edit routing-instances "RetailerInstance1" forwarding-options dhcp-relay]
user@host# edit authentication
```

9. (Optional) Configure a password that authenticates the username to the external authentication service.

See *Example-Configuring DHCP with External Authentication Server*.

10. (Optional) Configure optional features to create a unique username.

See *Creating Unique Usernames for DHCP Clients*.

11. Specify the default dynamic profile that you want to attach to DHCP subscriber for this retailer.

```
[edit routing-instances "RetailerInstance1" forwarding-options dhcp-relay]
user@host# set dynamic-profile Subscriber_Profile_Retail1
```

12. Specify any overrides for the default DHCP Relay configuration.

See *Overriding the Default DHCP Relay Configuration Settings*.

13. Configure a named server group for the retailer.

```
[edit routing-instances "RetailerInstance1" forwarding-options dhcp-relay]
user@host# edit server-group Retailer1_Group
```

14. Specify the DHCP server address for the retailer group.

```
[edit routing-instances "RetailerInstance1" forwarding-options dhcp-relay server-group "Retailer1_Group"]
user@host# set 10.10.100.1
```

15. Specify the retailer group as the active server group for this routing instance.

```
[edit routing-instances "RetailerInstance1" forwarding-options dhcp-relay]
user@host# set active-server-group Retailer1_Group
```

16. Configure a group you can use to define the retailer dynamic profile and DHCP access interface.

```
[edit routing-instances "RetailerInstance1" forwarding-options dhcp-relay]
user@host# edit group Retailer1_Group
```

17. Specify the dynamic profile that the retailer DHCP subscribers use.

```
[edit routing-instances "RetailerInstance1" forwarding-options dhcp-relay group "Retailer1_Group"]
user@host# set dynamic-profile Subscriber_Profile_Retailer1
```

18. Specify the retailer interface that the retailer DHCP subscribers use.

```
[edit routing-instances "RetailerInstance1" forwarding-options dhcp-relay group "Retailer1_Group"]
user@host# set interface ge-2/3/0.2
```

19. (Optional) Configure any passwords that authenticate the username to the external authentication service for the retailer groups that you created.

See Example-Configuring DHCP with External Authentication Server.

20. (Optional) Configure any unique username values for the retailer groups that you created.

See Creating Unique Usernames for DHCP Clients.

21. (Optional) Specify any overrides for any of the DHCP Relay group configurations that you created.

See Overriding the Default DHCP Relay Configuration Settings.

22. Repeat this procedure for other retailers.

RELATED DOCUMENTATION

Configure Default Forwarding Options for the DHCPv4 Wholesale Network Solution

You can use DHCP Relay, DHCP Relay Proxy, or DHCP Local Server configuration in a DHCP wholesale network. DHCP configuration is defined at the **[edit forwarding-options]** hierarchy level.

NOTE: The configuration for this wholesale solution uses DHCP Relay.

To configure DHCPv4 Relay forwarding options:

1. Access the **[edit forwarding-options dhcp-relay]** hierarchy.

```
[edit]
user@host# edit forwarding-options dhcp-relay
```

2. Specify that you want to configure authentication options and use external AAA authentication services.

```
[edit forwarding-options dhcp-relay]
user@host# edit authentication
```

3. (Optional) Configure a password that authenticates the username to the external authentication service.

See *Example-Configuring DHCP with External Authentication Server*.

4. (Optional) Configure optional features to create a unique username.

See *Creating Unique Usernames for DHCP Clients*.

5. Specify the default dynamic profile that you want to attach to all DHCP subscriber that access the router.

```
[edit forwarding-options dhcp-relay]
user@host# set dynamic-profile Wholesaler_Profile
```

6. Specify any overrides for the default DHCP Relay configuration.

See *Overriding the Default DHCP Relay Configuration Settings*.

7. Configure a named server group for default (wholesaler) DHCP server access.

```
[edit forwarding-options dhcp-relay]
user@host# edit server-group Wholesaler_Group
```

8. Specify the DHCP server address for the default (wholesale) group.

```
[edit forwarding-options dhcp-relay server-group "Wholesaler_Group"]
user@host# set 192.168.100.1
```

9. Specify the default (wholesale) group as the active server group.

```
[edit forwarding-options dhcp-relay]
user@host# set active-server-group Wholesaler_Group
```

10. Configure a group you can use to define the wholesale DHCP access interface.

```
[edit forwarding-options dhcp-relay]
user@host# edit group Wholesaler_Group
```

11. Specify the default (wholesale) interface that all DHCP subscribers use when first accessing the router.

```
[edit forwarding-options dhcp-relay group "Wholesaler_Group"]
user@host# set interface ge-2/3/0.1
```

12. Configure a group you can use to define a retail DHCP interface.

```
[edit forwarding-options dhcp-relay]
user@host# edit group Retailer1_Group
```

13. Specify the logical interface the DHCP subscribers use once redirected.

```
[edit forwarding-options dhcp-relay group "Retailer1_Group"]
user@host# set interface ge-2/3/0.2
```

14. Repeat steps 12 and 13 for other retailer groups.

In this solution example, you configure another group name of "Retailer2_Group" and specify **ge-2/3/0.3** for the logical interface.

15. (Optional) Configure any passwords that authenticate the username to the external authentication service for any of the groups that you created.

See *Example-Configuring DHCP with External Authentication Server*.

16. (Optional) Configure optional features to create a unique username for any of the groups that you created.

See *Creating Unique Usernames for DHCP Clients*.

17. (Optional) Specify any overrides for any of the DHCP Relay group configurations that you created.

See *Overriding the Default DHCP Relay Configuration Settings*.

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview

DHCP Relay Proxy Overview

Example-Configuring DHCP with External Authentication Server

Creating Unique Usernames for DHCP Clients

Overriding the Default DHCP Relay Configuration Settings

Example: Wholesaler Dynamic Profile for a DHCPv4 Wholesale Network

This example specifies a dynamic profile name of *Wholesaler_Profile*, uses dynamic IP demux interfaces, and references the predefined input firewall filter.

```
dynamic-profiles {
  Wholesaler_Profile {
    interfaces {
      demux0 {
        unit "$junos-interface-unit" {
          demux-options {
            underlying-interface "$junos-underlying-interface";
          }
        }
        family inet {
          demux-source {
            $junos-subscriber-ip-address;
          }
        }
        filter {
```

```

        input "$junos-input-filter";
    }
    unnumbered-address "$junos-loopback-interface" preferred-source-address
        $junos-preferred-source-address;
    }
}
}
}
}

```

RELATED DOCUMENTATION

[Configuring Dynamic Profiles for the DHCPv4 Layer 3 Wholesale Network Solution](#) | 20

Example: Retailer Dynamic Profile for a DHCPv4 Wholesale Network

```

dynamic-profiles {
  Subscriber_Profile_Retailer1 {
    routing-instances {
      "$junos-routing-instance" {
        interface "$junos-interface-name";
      }
    }
    interfaces {
      demux0 {
        unit "$junos-interface-unit" {
          demux-options {
            underlying-interface "$junos-underlying-interface";
          }
          family inet {
            demux-source {
              "$junos-subscriber-ip-address";
            }
            unnumbered-address "$junos-loopback-interface" preferred-source-address
              "$junos-preferred-source-address";
          }
        }
      }
    }
  }
}

```

```
}
```

RELATED DOCUMENTATION

[Configuring Dynamic Profiles for the DHCPv4 Layer 3 Wholesale Network Solution](#) | 20

Example: Default Forwarding Options Configuration for the DHCPv4 Wholesale Network

```
forwarding-options {
  dhcp-relay {
    traceoptions {
      file size 1g;
      inactive: flag all;
    }
    authentication {
      password $ABC123;
      username-include {
        user-prefix WholesaleNetwork;
      }
    }
  }
  dynamic-profile Wholesaler_Profile;
  overrides {
    always-write-giaddr;
    always-write-option-82;
    layer2-unicast-replies;
    trust-option-82;
    client-discover-match;
  }
  server-group {
    Wholesaler-Server-Group {
      192.168.100.1;
    }
  }
  active-server-group Wholesaler-Server Group;
  group Wholesaler-Group {
    authentication {
      password $ABC123;
      username-include {
```

```

        user-prefix WholesaleNetwork;
    }
}
interface ge-2/3/0.1;
}
group Retailer1-Group {
    authentication {
        password $ABC123$ABC123;
        username-include {
            user-prefix WholesaleNetwork_Retailer1;
        }
    }
    interface ge-2/3/0.2;
}
group Retailer2-Group {
    authentication {
        password $ABC123$ABC123$ABC123;
        username-include {
            user-prefix WholesaleNetwork_Retailer1;
        }
    }
    interface ge-2/3/0.3;
}
}
}

```

RELATED DOCUMENTATION

| [Configure Default Forwarding Options for the DHCPv4 Wholesale Network Solution](#) | 27

Example: Retailer Routing Instances for a DHCPv4 Wholesale Network

```

routing-instances {
    Retailer_Instance1 {
        instance-type vrf;
        access-profile Retailer_Access1;
        interface ge-11/1/9.10;
        interface ge-11/1/10.100;
        interface lo0.1;
    }
}

```

```

route-distinguisher 1:1;
forwarding-options {
  dhcp-relay {
    authentication {
      password $ABC123$ABC123;
      username-include {
        user-prefix WholesaleNetwork_Retailer1;
      }
    }
  }
  dynamic-profile Subscriber_Profile_Retailer1;
  overrides {
    always-write-giaddr;
    always-write-option-82;
    layer2-unicast-replies;
    trust-option-82;
    client-discover-match;
  }
  server-group {
    Retailer1-Server-Group {
      10.10.100.1;
    }
  }
  active-server-group Retailer1-Server-Group;
  group Retailer1-Group {
    authentication {
      password $ABC123$ABC123;
      username-include {
        user-prefix WholesaleNetwork_Retailer1;
      }
    }
    dynamic-profile Subscriber_Profile_Retailer1;
    overrides {
      always-write-giaddr;
      trust-option-82;
      client-discover-match;
    }
    interface ge-2/3/0.2;
  }
}
}
}
Retailer_Instance2 {
  instance-type vrf;
  access-profile Retailer_Access2;
}

```

```

interface ge-7/1/9.10;
interface ge-7/1/9.100;
interface lo0.2;
route-distinguisher 2:2;
forwarding-options {
  dhcp-relay {
    authentication {
      password $ABC123$ABC123$ABC123;
      username-include {
        user-prefix WholesaleNetwork_Retailer2;
      }
    }
    dynamic-profile Subscriber_Profile_Retailer2;
    overrides {
      always-write-giaddr;
      trust-option-82;
      client-discover-match;
    }
    server-group {
      Retailer2-Group {
        10.20.200.1;
      }
    }
    active-server-group Retailer2-Group;
    group Retailer2-Group {
      authentication {
        password $ABC123$ABC123$ABC123;
        username-include {
          user-prefix psswd2;
        }
      }
      dynamic-profile Subscriber_Profile_Retailer2;
      overrides {
        always-write-giaddr;
        trust-option-82;
        client-discover-match;
      }
      interface ge-2/3/0.3;
    }
  }
}
}
}
}

```

RELATED DOCUMENTATION

| [Configuring Separate Routing Instances for DHCPv4 Service Retailers](#) | 23

Configuring DHCPv6 Layer 3 Wholesale Networks

IN THIS CHAPTER

- Broadband Subscriber Management DHCPv6 Layer 3 Wholesale Topology and Configuration Elements | 36
- DHCPv6 Layer 3 Wholesale Network Topology Overview | 38
- Configuring Loopback Interfaces for the DHCPv6 Layer 3 Wholesale Solution | 39
- Configuring VLANs for the DHCPv6 Layer 3 Wholesale Network Solution | 40
- Configuring Access Components for the DHCP Layer 3 Wholesale Network Solution | 44
- Configuring Dynamic Profiles for the DHCPv6 Layer 3 Wholesale Network Solution | 47
- Configuring Separate Routing Instances for DHCPv6 Service Retailers | 51
- Configuring Address Server Elements for the DHCPv6 Layer 3 Wholesale Solution | 52
- Example: Retailer Dynamic Profile for a DHCPv6 Wholesale Network | 56
- Example: Retailer Routing Instances for a DHCPv6 Wholesale Network | 57
- Example: DHCPv6 Address Assignment Pool That Provides Full 128-bit IPV6 Addresses for a DHCPv6 Wholesale Network | 58
- Example: DHCPv6 Address Assignment Pool That Provides 74-bit IPV6 Prefixes for a DHCPv6 Wholesale Network | 59
- Example: Extended DHCPv6 Local Server for a DHCPv6 Wholesale Network | 59

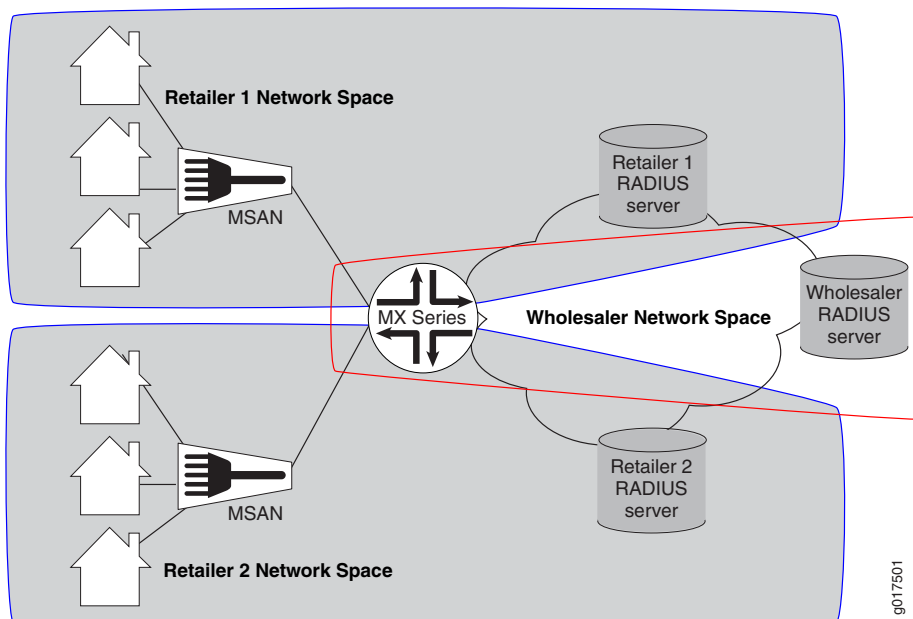
Broadband Subscriber Management DHCPv6 Layer 3 Wholesale Topology and Configuration Elements

The network topology for the subscriber management DHCPv6 Layer 3 wholesale solution includes configuring separate routing instances for individual retailers that use a portion of the router. This solution uses a DHCPv6 local server configuration.

NOTE: Only DHCPv6 local server is currently supported for DHCPv6 Layer 3 wholesale configuration.

To explain the concept, but to limit complexity, this solution provides a configuration with one wholesaler and only two retailers. [Figure 3 on page 37](#) illustrates a basic Layer 3 wholesale topology model from which you can expand.

Figure 3: Basic Subscriber Management DHCPv6 Layer 3 Wholesale Solution Topology



A DHCPv6 Layer 3 wholesale network solution can use various combinations of the following configuration elements:

- Subscriber network VLAN configuration
- DHCPv6 configuration (local server only)
- RADIUS server access configuration
- Dynamic profile configuration for default (wholesaler) access
- Dynamic profile configuration for retailer access (following subscriber redirection, if applicable)
- Routing instance configuration for individual retailers
- Group configuration and forwarding options for the network
- Core network configuration

RELATED DOCUMENTATION

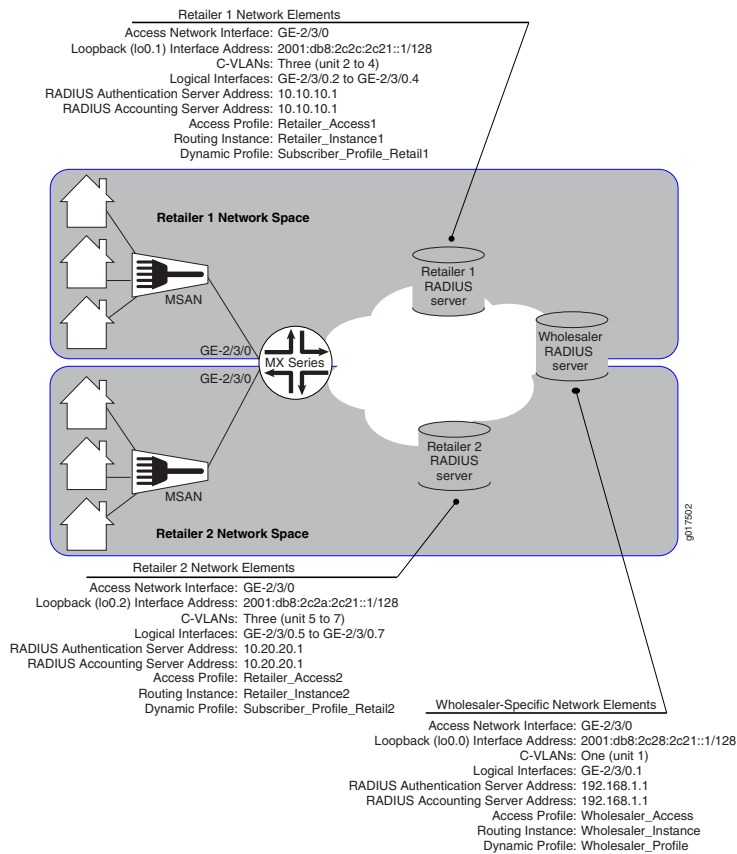
[Layer 2 and Layer 3 Wholesale Overview | 2](#)

[DHCPv6 Layer 3 Wholesale Network Topology Overview | 38](#)

DHCPv6 Layer 3 Wholesale Network Topology Overview

This configuration explains how to configure a simple DHCPv6 Layer 3 wholesale subscriber access network. This solution incorporates two retailers sharing resources on a wholesaler router. [Figure 4 on page 38](#) provides the reference topology for this configuration example.

Figure 4: DHCPv6 Layer 3 Wholesale Network Reference Topology



RELATED DOCUMENTATION

[Layer 2 and Layer 3 Wholesale Overview | 2](#)

[Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Topology and Configuration Elements | 8](#)

Configuring Loopback Interfaces for the DHCPv6 Layer 3 Wholesale Solution

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

To configure loopback interfaces:

1. Edit the loopback interface.

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

2. Edit the unit for the loopback interface that you want to use for the wholesaler.

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

3. Edit the loopback interface family that belongs to the wholesaler.

```
[edit interfaces lo0 unit 0]
user@host# edit family inet6
```

4. Specify the wholesale loopback interface address.

```
[edit interfaces lo0 unit 0]
user@host# set address 2001:db8:2c28:2c21::1/128
```

5. Edit the unit for a retail loopback interface.

```
[edit interfaces lo0]
user@host# edit unit 1
```

6. Edit the retail loopback interface family.

```
[edit interfaces lo0 unit 1]
user@host# edit family inet6
```

7. Specify the retail loopback interface address.

```
[edit interfaces lo0 unit 1]
```

```
user@host# set address 2001:db8:2c2c:2c21::1/128
```

8. Repeat steps 5 through 7 for additional retailers, making sure to use unique unit and address values for each retailer loopback interface.

Configuring VLANs for the DHCPv6 Layer 3 Wholesale Network Solution

IN THIS SECTION

- [Configuring Static Customer VLANs for the DHCPv6 Layer 3 Wholesale Network Solution | 40](#)
- [Configuring Dynamic Customer VLANs for the DHCPv6 Layer 3 Wholesale Network Solution | 41](#)

You can configure either static or dynamic customer VLANs for use in the DHCPv6 wholesale network solution.

Configuring Static Customer VLANs for the DHCPv6 Layer 3 Wholesale Network Solution

In this example configuration, the access interface (**ge-2/3/0**) connects to a device (that is, a DSLAM) on the access side of the network. You can define static VLANs for use by access network subscribers.

To configure the static VLANs:

1. Edit the access side interface.

```
[edit]
user@host# edit interfaces ge-2/3/0
```

2. Specify the use of stacked VLAN tagging.

```
[edit interfaces ge-2/3/0]
user@host# set stacked-vlan-tagging
```

3. Edit the interface unit for the first VLAN.

```
[edit interfaces ge-2/3/0]
user@host# edit unit 1
```

4. Define the VLAN tags for the first VLAN.

```
[edit interfaces ge-2/3/0 unit 1]
user@host# set vlan-tags outer 3 inner 1
```

5. Specify that you want to create IPv6 demux interfaces.

```
[edit interfaces ge-2/3/0 unit 1]
user@host# set demux-source inet6
```

6. Edit the family for the first VLAN.

```
[edit interfaces ge-2/3/0 unit 1]
user@host# edit family inet6
```

7. (Optional) Define the unnumbered address and the preferred source address for the first VLAN.

```
[edit interfaces ge-2/3/0 unit 1 family inet6]
user@host# set unnumbered-address lo0.1 preferred-source-address 2001:db8:2c28:2c21::1/128
```

8. Repeat steps 2 through 7 for additional VLAN interface units.

Configuring Dynamic Customer VLANs for the DHCPv6 Layer 3 Wholesale Network Solution

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# edit unit $junos-interface-unit
```

- d. Specify that you want to create IPv6 demux interfaces.

```
[edit dynamic-profiles VLAN-PROF interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit"]
user@host# set demux-source inet6
```

- e. 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 **[interfaces]** hierarchy level.

- f. 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 **[interfaces]** hierarchy level.

- g. Access the family type.

```
[edit dynamic-profiles VLAN-PROF interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit"]
user@host# edit family inet6
```

- h. (Optional) Specify the unnumbered address and preferred source address.

```
[edit dynamic-profiles VLAN-PROF interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit"
family inet6]
user@host# set unnumbered-address lo.0 preferred-source-address 2001:db8:2c28:2c21::1/128
```

2. Associate the dynamic profile with the interface on which you want the VLANs created.

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

```
[edit interfaces]
user@host# edit interfaces ge-2/3/0
```

- b. Specify the use of stacked VLAN tagging.

```
[edit interfaces ge-2/3/0]
user@host# set stacked-vlan-tagging
```

- c. Specify that you want to automatically configure VLAN interfaces.

```
[edit interfaces ge-2/3/0]
user@host# edit auto-configure
```

- d. Specify that you want to configure stacked VLANs.

```
[edit interfaces ge-2/3/0 auto-configure]
user@host# edit stacked-vlan-ranges
```

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

```
[edit interfaces ge-2/3/0 auto-configure stacked-vlan-ranges]
user@host# set dynamic-profile VLAN-PROF
```

- f. Repeat steps a through e for any other interfaces that you want to use for creating VLANs.

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

```
[edit interfaces ge-2/3/0 auto-configure stacked-vlan-ranges dynamic-profile VLAN-PROF]
user@host# set accept inet6
```

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. 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 **1-3** (enabling a range from 1 through 3 for the inner stacked VLAN ID).

```
[edit interfaces ge-0/0/0 auto-configure stacked-vlan-ranges dynamic-profile VLAN-PROF]
```

```
user@host# set stacked-vlan-ranges 3-3,1-3
```

Configuring Access Components for the DHCP Layer 3 Wholesale Network Solution

IN THIS SECTION

- [Configuring RADIUS Server Access | 44](#)
- [Configuring a DHCP Wholesaler Access Profile | 45](#)
- [Configuring DHCP Retailer Access Profiles | 46](#)

When configuring a wholesale network, you must configure several components globally. This configuration provides access to RADIUS servers that you want the wholesaler and any configured retailers to use globally. The access configuration includes the following general steps:

Configuring RADIUS Server Access

You can globally define any RADIUS servers in your network that either the wholesaler access profile or retailer access profile can use. After you define the global RADIUS servers, you can specify specific RADIUS servers within individual access profiles.

To define RADIUS servers for profile access:

1. Access the **[edit access radius-server]** hierarchy level.

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

2. Specify the address and secret for any RADIUS servers in the network.

```
[edit access radius-server]
user@host# set 192.168.10.1 secret $ABC123$ABC123$ABC123
user@host# set 10.10.10.1 secret $ABC123$ABC123
```


SEE ALSO

| *RADIUS Servers and Parameters for Subscriber Access*

Configuring a DHCP Wholesaler Access Profile

You must define the network and interface over which you want subscribers to initially access the network with a wholesale access profile. When a subscriber attempts to access the network, the access profile provides initial access information including authentication and accounting values that the router uses for the accessing subscriber.

To define a wholesale access profile:

1. Create the wholesale access profile.

```
[edit]
user@host# edit access-profile Wholesaler_Access
```

2. Specify the authentication methods for the profile and the order in which they are used.

```
[edit access profile Wholesaler1]
user@host# set authentication-order radius password
```

3. Specify that you want to configure RADIUS support.

```
[edit access profile Wholesaler1]
user@host# edit radius
```

4. Specify the IP address of the RADIUS server used for authentication.

```
[edit access profile Wholesaler1 radius]
user@host# set authentication-server 192.168.10.1
```

5. Specify the IP address of the RADIUS server used for accounting.

```
[edit access profile Wholesaler1 radius]
user@host# set accounting-server 192.168.10.1
```

6. Configure any desired options for the RADIUS server.

See *RADIUS Servers and Parameters for Subscriber Access*.

7. Configure subscriber accounting (RADIUS accounting).

See *Configuring Per-Subscriber Session Accounting*.

SEE ALSO

Configuring Authentication and Accounting Parameters for Subscriber Access

Specifying the Authentication and Accounting Methods for Subscriber Access

RADIUS Servers and Parameters for Subscriber Access

Configuring Per-Subscriber Session Accounting

Configuring DHCP Retailer Access Profiles

In this solution, subscribers are redirected to a networking space used by a specific retailer and defined by a unique routing instance. This method requires that you define the network and interface over which you want subscribers to access the network after being redirected by the wholesale access profile.

To define a retailer access profile:

1. Create the retailer access profile.

```
[edit]
user@host# edit access-profile Retailer_Access1
```

2. Specify the authentication methods for the profile and the order in which they are used.

```
[edit access profile Retailer1]
user@host# set authentication-order radius password
```

3. Specify that you want to configure RADIUS support.

```
[edit access profile Retailer1]
user@host# edit radius
```

4. Specify the IP address of the RADIUS server used for authentication.

```
[edit access profile Retailer1 radius]
```

```
user@host# set authentication-server 10.10.10.1
```

5. Specify the IP address of the RADIUS server used for accounting.

```
[edit access profile Retailer1 radius]
user@host# set accounting-server 10.10.10.1
```

6. Configure any desired options for the RADIUS server.

See *RADIUS Servers and Parameters for Subscriber Access*.

7. Configure subscriber accounting (RADIUS accounting).

See *Configuring Per-Subscriber Session Accounting*.

SEE ALSO

Configuring Authentication and Accounting Parameters for Subscriber Access

Specifying the Authentication and Accounting Methods for Subscriber Access

RADIUS Servers and Parameters for Subscriber Access

Configuring Per-Subscriber Session Accounting

Configuring Dynamic Profiles for the DHCPv6 Layer 3 Wholesale Network Solution

IN THIS SECTION

- [Configuring a Wholesale Dynamic Profile for use in the DHCPv6 Solution | 48](#)
- [Configuring a Dynamic Profile for use by Each Retailer in the DHCPv6 Solution | 49](#)

A dynamic profile is a set of characteristics, defined in a type of template, that you can use to provide services for broadband applications. These services are assigned dynamically to interfaces as they access the network. When configuring dynamic profiles for the DHCPv6 Layer 3 wholesale network, you can

choose to configure one dynamic profile to address all incoming subscribers or you can configure individual dynamic profiles for use by the different network management groups (that is, the wholesaler and any retailers). In fact, you can create multiple dynamic profiles that you can use to roll out different services and selectively apply those dynamic profiles to different subscriber groups as necessary.

In this solution example, one dynamic profile is created for use by the wholesaler when subscribers initially access the network. Other dynamic profiles are created for the subscribers for each individual retailer to use after they are redirected to that retailer network space.

Configuring a Wholesale Dynamic Profile for use in the DHCPv6 Solution

You can configure a basic access profile to initially manage subscribers that access the network.

To configure a dynamic profile for use by the wholesaler:

1. Create a wholesale dynamic profile.

```
[edit]
user@host# edit dynamic-profiles Wholesaler_Profile
```

2. Specify that you want to configure the **demux0** interface in the dynamic profile.

```
[edit dynamic-profiles Wholesaler_Profile]
user@host# edit interfaces demux0
```

3. Configure the unit for the **demux0** interface.

- a. Configure the variable for the unit number of the **demux0** interface.

The variable is dynamically replaced with the unit number that DHCP supplies when the subscriber logs in.

```
[edit dynamic-profiles Wholesaler_Profile demux0]
user@host# edit unit $junos-interface-unit
```

- b. Configure the variable for the underlying interface of the demux interfaces and specify the **\$junos-underlying-interface** variable.

The variable is dynamically replaced with the underlying interface that DHCP supplies when the subscriber logs in.

```
[edit dynamic-profiles Wholesaler_Profile interfaces demux0 unit "$junos-interface-unit"]
```

```
user@host# set demux-options underlying-interface $junos-underlying-interface
```

4. Configure the family for the demux interfaces.

- a. Specify that you want to configure the family.

```
[edit dynamic-profiles Wholesaler_Profile interfaces demux0 unit "$junos-interface-unit"]
user@host# edit family inet6
```

- b. Configure the unnumbered address for the family.

```
[edit dynamic-profiles Wholesaler_Profile demux0 unit "$junos-interface-unit" family inet6]
user@host# set unnumbered-address lo0.0
```

- c. Configure the variable for the IPv6 address of the demux interface.

The variable is dynamically replaced with the IPv6 address that DHCP supplies when the subscriber logs in.

```
[edit dynamic-profiles Wholesaler_Profile interfaces demux0 unit "$junos-interface-unit"]
user@host# set demux-source $junos-subscriber-ipv6-address
```

Configuring a Dynamic Profile for use by Each Retailer in the DHCPv6 Solution

To configure a dynamic profile for use with retailer access:

1. Create a retail dynamic profile.

```
[edit]
user@host# edit dynamic-profiles Subscriber_Profile_Retail1
```

2. Define the dynamic routing instance variable in the dynamic profile.

```
[edit dynamic-profiles Subscriber_Profile_Retail1]
user@host# edit routing-instances $junos-routing-instance
```

3. Set the dynamic interface variable for the dynamic routing instance.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 routing-instances "$junos-routing-instance"]
user@host# set interface $junos-interface-name
```

4. Specify that you want to configure the **demux0** interface in the dynamic profile.

```
[edit dynamic-profiles Subscriber_Profile_Retail1]
user@host# edit interfaces demux0
```

5. Configure the unit for the **demux0** interface.

- a. Configure the variable for the unit number of the **demux0** interface.

The variable is dynamically replaced with the unit number that DHCP supplies when the subscriber logs in.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 demux0]
user@host# edit unit $junos-interface-unit
```

- b. Configure the variable for the underlying interface of the demux interfaces and specify the **\$junos-underlying-interface** variable.

The variable is dynamically replaced with the underlying interface that DHCP supplies when the subscriber logs in.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces demux0 unit "$junos-interface-unit"]
user@host# set demux-options underlying-interface $junos-underlying-interface
```

6. Configure the family for the demux interfaces.

- a. Specify that you want to configure the family.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces demux0 unit "$junos-interface-unit"]
user@host# edit family inet6
```

- b. Configure the unnumbered address and preferred source address for the family.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 demux0 unit "$junos-interface-unit" family inet6]
user@host# set unnumbered-address $junos-loopback-interface preferred-source-address
$junos-preferred-source-address
```

- c. Configure the variable that identifies the demux interface on the logical interface.

The variable is dynamically replaced with the IPv6 address that DHCP supplies when the subscriber logs in.

```
[edit dynamic-profiles business-profile interfaces demu0 unit "$junos-interface-unit"]
user@host# set demux-source $junos-subscriber-ipv6-address
```

Configuring Separate Routing Instances for DHCPv6 Service Retailers

As the owner of the system, the wholesaler typically uses the default routing instance. You must create separate routing instances for each individual retailer to keep routing information for individual retailers separate and to define any servers and forwarding options specific to each retailer.

To define a retailer routing instance:

1. Create the retailer routing instance.

```
[edit]
user@host# edit routing-instances Retailer_Instance1
```

2. Specify the routing instance type for the retailer.

```
[edit routing-instances "Retailer_Instance1"]
user@host# set instance-type vrf
```

3. Specify the access profile that you want the routing instance to use.

```
[edit routing-instances "Retailer_Instance1"]
user@host# set access-profile Retailer_Access1
```

4. Specify the interface that faces the Retailer1 RADIUS server.

```
[edit routing-instances "Retailer_Instance1"]
user@host# set interface ge-11/1/9.10
```

5. Specify the loopback interface unit for this routing instance.

```
[edit routing-instances "RetailerInstance1"]
user@host# set interface lo0.1
```

NOTE: Loopback interfaces must be unique for each routing instance.

6. Repeat this procedure for other retailers.

RELATED DOCUMENTATION

Configuring Address Server Elements for the DHCPv6 Layer 3 Wholesale Solution

IN THIS SECTION

- [Configuring a DHCPv6 Address Assignment Pool | 52](#)
- [Configuring Extended DHCPv6 Local Server | 54](#)

Configuring a DHCPv6 Address Assignment Pool

Address assignment pools enable you to specify groups of IPv6 addresses that different client applications can share. In this configuration, the extended DHCPv6 local server configuration uses the address pool to provide addresses to subscribers that are accessing the network. You must create separate address assignment pools for each retailer routing instance.

You can create address assignment pools that provide full 128 bit IPv6 addresses or pools that provide prefixes of a specified length.

To configure an address assignment pool that provides full 128 -bit IPv6 addresses:

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 inet6
```

3. Define the IPv6 network prefix.

```
[edit access address-pool AddressPool_1 family inet6]
user@host# set prefix 2001:db8:2121::0/64
```

4. Define a named address range for the pool of IPv6 addresses.

```
[edit access address-assignment pool AddressPool_1 family inet6]
user@host# set range Range1 low 2001:db8:2121::a/128
user@host# set range Range1 high 2001:db8:2121::7ffe/128
```

5. (Optional) Edit the family DHCP attributes.

```
[edit access address-assignment pool AddressPool_1 family inet6]
user@host# edit dhcp-attributes
```

6. (Optional) Set the maximum lease time.

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

7. (Optional) Set the grace period.

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

To configure an address assignment pool that provides shorter, 74-bit IPv6 prefixes:

1. Create and name an address assignment pool.

```
[edit]
user@host# edit access address-assignment pool AddressPool_2
```

2. Edit the address pool family.

```
[edit access address-assignment pool AddressPool_2]
user@host# edit family inet6
```

3. Define the IPv6 network prefix.

```
[edit access address-pool AddressPool_2 family inet6]
user@host# set prefix 2001:db8:2222::0/64
```

4. Define a named address range limit for the pool of IPv6 addresses.

```
[edit access address-assignment pool AddressPool_2 family inet6]
user@host# set range BitLimit prefix-length 74
```

5. (Optional) Edit the family DHCP attributes.

```
[edit access address-assignment pool AddressPool_2 family inet6]
user@host# edit dhcp-attributes
```

6. (Optional) Set the maximum lease time.

```
[edit access address-assignment pool AddressPool_2 family inet dhcp-attributes]
user@host# set maximum-lease-time 3600
```

7. (Optional) Set the grace period.

```
[edit access address-assignment pool AddressPool_2 family inet dhcp-attributes]
user@host# set grace-period 60
```

Configuring Extended DHCPv6 Local Server

You can enable the MX Series router to function as an extended DHCPv6 local server. The extended DHCPv6 local server provides IPv6 addresses and other configuration information to a subscriber logging

into the network. You must configure extended DHCPv6 local server for the wholesaler (default) routing instance and also for each retailer routing instance.

To configure the DHCPv6 local server:

1. Edit the routing system services.

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

2. Edit the DHCPv6 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 $ABC123
```

5. (Optional) Edit the values you want included with the username.

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

6. (Optional) Set the values you want included with the username.

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

7. Access the DHCPv6-specific service configuration.

```
[edit system services dhcp-local-server]
user@host# edit dhcpv6
```

8. Create and name a DHCPv6 local server group.

```
[edit system services dhcp-local-server dhcpv6]
user@host# edit group dhcp-ls-group
```

9. Specify a dynamic profile that you want the DHCPv6 local server group to use.

```
[edit system services dhcp-local-server dhcpv6 group dhcp-ls-group]
user@host# set dynamic-profile Wholesaler_Profile
```

10. Assign interfaces to the group.

```
[edit system services dhcp-local-server dhcpv6 group dhcp-ls-group]
user@host# set interface ge-1/3/0.1 upto ge-1/3/0.5
```

11. Edit the DHCPv6 local server trace options.

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

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

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

13. Specify the DHCPv6 local server message operations that you want saved in the log file.

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

RELATED DOCUMENTATION

Address-Assignment Pools Overview

DHCPv6 Local Server Overview

Example: Retailer Dynamic Profile for a DHCPv6 Wholesale Network

```
dynamic-profiles {
  Subscriber_Profile_Retailer1 {
```

```

routing-instances {
  "$junos-routing-instance" {
    interface "$junos-interface-name";
  }
}
interfaces {
  demux0 {
    unit "$junos-interface-unit" {
      demux-options {
        underlying-interface "$junos-underlying-interface";
      }
      family inet6 {
        demux-source {
          "$junos-subscriber-ip-address";
        }
        unnumbered-address "$junos-loopback-interface" preferred-source-address
          "$junos-preferred-source-address";
      }
    }
  }
}
}
}

```

RELATED DOCUMENTATION

[Configuring Dynamic Profiles for the DHCPv6 Layer 3 Wholesale Network Solution](#) | 47

Example: Retailer Routing Instances for a DHCPv6 Wholesale Network

```

routing-instances {
  Retailer_Instance1 {
    instance-type vrf;
    access-profile Retailer_Access1;
    interface ge-11/1/9.10;
    interface lo0.1;
    route-distinguisher 1:1;
  }
  Retailer_Instance2 {
    instance-type vrf;

```

```

    access-profile Retailer_Access2;
    interface ge-7/1/9.10;
    interface lo0.2;
  }
}

```

RELATED DOCUMENTATION

[Configuring Separate Routing Instances for DHCPv6 Service Retailers](#) | 51

Example: DHCPv6 Address Assignment Pool That Provides Full 128-bit IPV6 Addresses for a DHCPv6 Wholesale Network

```

access {
  address-assignment {
    pool AddressPool_1 {
      family inet6 {
        prefix 2001:db8:2121::0/64;
        range Range1 {
          low 2001:db8:2121::a/128;
          high 2001:db8:2121::7ffe/128;
        }
        dhcp-attributes {
          maximum-lease-time 3600;
          grace-period 60;
        }
      }
    }
  }
}

```

RELATED DOCUMENTATION

[DHCPv6 Address Assignment Pool That Provides Full 128-bit IPV6 Addresses for a DHCPv6 Wholesale Network](#)

Example: DHCPv6 Address Assignment Pool That Provides 74-bit IPV6 Prefixes for a DHCPv6 Wholesale Network

```
access {
  address-assignment {
    pool AddressPool_2 {
      family inet6 {
        prefix 2001:db8:2222::0/64;
        range BitLimit prefix-length 74;
        dhcp-attributes {
          maximum-lease-time 3600;
          grace-period 60;
        }
      }
    }
  }
}
```

RELATED DOCUMENTATION

[Configuring Address Server Elements for the DHCPv6 Layer 3 Wholesale Solution](#) | 52

Example: Extended DHCPv6 Local Server for a DHCPv6 Wholesale Network

```
system {
  services {
    dhcp-local-server {
      traceoptions {
        file dhcp-server-msgs.log;
        flag all;
      }
    }
    dhcpv6 {
      group dhcp-ls-group {
        dynamic-profile Wholesaler_Profile;
        interface ge-1/3/0.1 {
          upto ge-1/3/0.5;
        }
      }
    }
  }
}
```

```
    }  
  }  
}  
pool-match-order {  
  ip-address-first;  
}  
authentication {  
  password $ABC123;  
  username-include {  
    domain-name example.com;  
    user-prefix user-defined-prefix;  
  }  
}  
}  
}  
}
```

RELATED DOCUMENTATION

| [Configuring Address Server Elements for the DHCPv6 Layer 3 Wholesale Solution](#) | 52

2

PART

Configuring PPPoE Layer 3 Wholesale Networks

[Subscriber Management PPPoE Wholesale Overview](#) | **62**

[Configuring PPPoE Layer 3 Wholesale Networks](#) | **66**

Subscriber Management PPPoE Wholesale Overview

IN THIS CHAPTER

- [Layer 2 and Layer 3 Wholesale Overview | 62](#)
- [PPPoE Layer 3 Wholesale Configuration Interface Support | 63](#)
- [Subscriber to Logical System and Routing Instance Relationship | 64](#)
- [RADIUS VSAs and Broadband Subscriber Management Wholesale Configuration Overview | 64](#)

Layer 2 and Layer 3 Wholesale Overview

In general, wholesaling broadband services allows service providers to resell broadband services and allows other providers to deploy their own services over the incumbent network. There are different methods to partitioning an access network for resale. The two most common approaches are based on either Layer 2 or Layer 3 information. Wholesale access is the process by which the access network provider (the *wholesaler*) partitions the access network into separately manageable and accountable subscriber segments for resale to other network providers (or *retailers*).

In a Layer 3 wholesale configuration, you partition the wholesaler access network at the network layer or the subscriber IP component by associating the IP component with a distinct Layer 3 domain. In a Layer 2 wholesale configuration, you partition the access network at the subscriber circuit or customer VLAN (C-VLAN) by backhauling the connection through the service provider backbone network to the subscribing retailer network where the access traffic can be managed at higher layers.

In a Junos OS Dynamic Host Configuration Protocol (DHCP) or Point-to-Point Protocol over Ethernet (PPPoE) subscriber access configuration, wholesale partitioning is accomplished through the use of logical systems and routing instances within the router. Logical systems offer a stricter partitioning of routing resources than routing instances. The purpose behind the use of logical systems is to distinctly partition the physical router into separate administrative domains. This partitioning enables multiple providers to administer the router simultaneously, with each provider having access only to the portions of the configuration relevant to their logical system. Junos OS supports up to 15 named logical systems in addition to the default logical system (that is, **inet.0**). Unless otherwise specified in configuration, all interfaces belong to the default logical system.

NOTE: This Junos OS release supports the use of only the default logical system. Partitioning currently occurs through the use of separate routing instances.

A logical system can have one or more routing instances. Typically used in Layer 3 VPN scenarios, a routing instance does not have the same level of administrative separation as a logical system because it does not offer administrative isolation. However, the routing instance defines a distinct routing table, set of routing policies, and set of interfaces.

RELATED DOCUMENTATION

[Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Topology and Configuration Elements | 8](#)

[Broadband Subscriber Management PPPoE Layer 3 Wholesale Topology and Configuration Elements | 66](#)

[Broadband Subscriber Management Layer 2 Wholesale Topology and Configuration Elements | 89](#)

PPPoE Layer 3 Wholesale Configuration Interface Support

PPPoE Layer 3 wholesale requires the use of PPP interfaces. This means that you must specify the PPO interface when configuring Layer 3 wholesaling in a PPPoE network.

For general additional information about configuring PPPoE interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*.

RELATED DOCUMENTATION

Junos OS Network Interfaces Library for Routing Devices

Configuring a PPPoE Dynamic Profile.

Configuring Dynamic PPPoE Subscriber Interfaces.

Subscriber to Logical System and Routing Instance Relationship

As subscriber sessions are established, subscriber to logical system/routing instance memberships are established by the AAA framework configured for the default logical system. When configuring Layer 3 wholesaling, you typically configure global (wholesale) information within the default (master) logical system and default routing instance. Incoming subscribers must then be authenticated, but this authentication can be handled in one of two ways:

- Single (wholesaler only) authentication—Incoming subscribers are authenticated by the wholesaler RADIUS server. After authentication, the subscribers are assigned values specified by dynamic profiles (routing instances, interfaces, and any configuration values) specific to a particular retailer.
- Dual (wholesaler and retailer) authentication—Sometimes referred to as *double-dip authentication*. Incoming subscribers are initially authenticated by RADIUS using the wholesale configuration. Authenticated subscribers are then redirected to other routing instances associated with individual retailer network space. When you redirect subscribers, and those subscribers are to be authenticated by AAA servers owned by individual retailers, the subscribers must be authenticated again by the AAA servers before they are provided an address and any dynamic profile values are assigned. After reauthentication, however, the subscribers are managed normally using any values specific to the retailer routing instance to which they are assigned.

RELATED DOCUMENTATION

See *Routing Instances Overview*.

RADIUS VSAs and Broadband Subscriber Management Wholesale Configuration Overview

You can use RADIUS to assign various values through the use of dynamic variables within dynamic profiles. However, the configuration of at least one of the two VSAs described in [Table 3 on page 6](#) is required for a wholesale network to function.

Table 4: Required Juniper Networks VSAs for the Broadband Subscriber Management Wholesale Network Solution

Attribute Number	Attribute Name	Description	Value
26-1	LSRI-Name	Client logical system/routing instance membership name. Allowed only from RADIUS server for “default” logical system/routing instance membership.	string: logical system:routing instance
26-25	Redirect-LSRI-Name	Client logical system/routing instance membership name indicating to which logical system/routing instance membership the request is redirected for user authentication.	string: logical system:routing instance

Specifying the **\$junos-routing-instance** dynamic variable in a dynamic profile triggers a RADIUS access-accept response of either the LSRI-Name VSA or the Redirect-LSRI-Name VSA. Returning an LSRI-Name attribute in the access-accept response provides the logical system and routing instance in which the logical interface is to be created and the router updates the session database with the specified routing instance value. Returning a Redirect-LSRI-Name attribute in the access-accept response results in the router immediately sending a second access-request message (sometimes referred to as a *double-dip*) to the RADIUS server specified by the logical system:routing instance attribute specified by the Redirect-LSRI-Name VSA.

NOTE: Attributes returned as a result of a second access-request message to the logical system/routing instance membership specified by the Redirect-LSRI-Name VSA override any prior attributes returned by initial access-accept responses to the default logical system/routing instance membership.

RELATED DOCUMENTATION

Juniper Networks VSAs Supported by the AAA Service Framework.

Configuring PPPoE Layer 3 Wholesale Networks

IN THIS CHAPTER

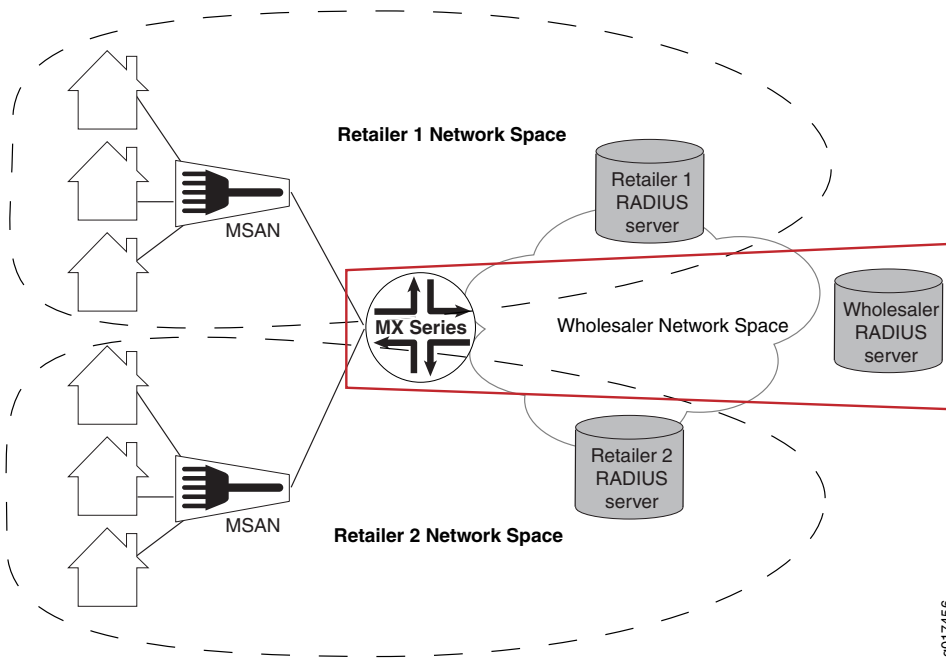
- Broadband Subscriber Management PPPoE Layer 3 Wholesale Topology and Configuration Elements | 66
- PPPoE Layer 3 Wholesale Network Topology Overview | 68
- Configuring Loopback Interfaces for the PPPoE Layer 3 Wholesale Solution | 69
- Configuring Static Customer VLANs for the PPPoE Layer 3 Wholesale Network Solution | 70
- Configuring Access Components for the PPPoE Wholesale Network Solution | 72
- Configuring Dynamic Profiles for the PPPoE Layer 3 Wholesale Network Solution | 75
- Configuring Separate Routing Instances for PPPoE Service Retailers | 78
- Example: Wholesaler Dynamic Profile for a PPPoE Wholesale Network | 79
- Example: Retailer Routing Instances for a PPPoE Wholesale Network | 80

Broadband Subscriber Management PPPoE Layer 3 Wholesale Topology and Configuration Elements

The network topology for the subscriber management PPPoE Layer 3 wholesale solution includes configuring separate routing instances for individual retailers that use a portion of the router.

To explain the concept, but to limit complexity, this solution provides a configuration with one wholesaler and only two retailers. [Figure 5 on page 67](#) illustrates a basic PPPoE Layer 3 wholesale topology model from which you can expand.

Figure 5: Basic Subscriber Management PPPoE Layer 3 Wholesale Solution Topology



When you are configuring a PPPoE Layer 3 wholesale network solution, the following configuration elements are required:

- Subscriber network VLAN configuration
- Addressing server or addressing server access configuration
- RADIUS server access configuration
- Dynamic profile configuration for default (wholesaler) access
- Routing instance configuration for individual retailers
- Group configuration and forwarding options for the network
- Core network configuration

This implementation of PPPoE Layer 3 wholesale supports the following:

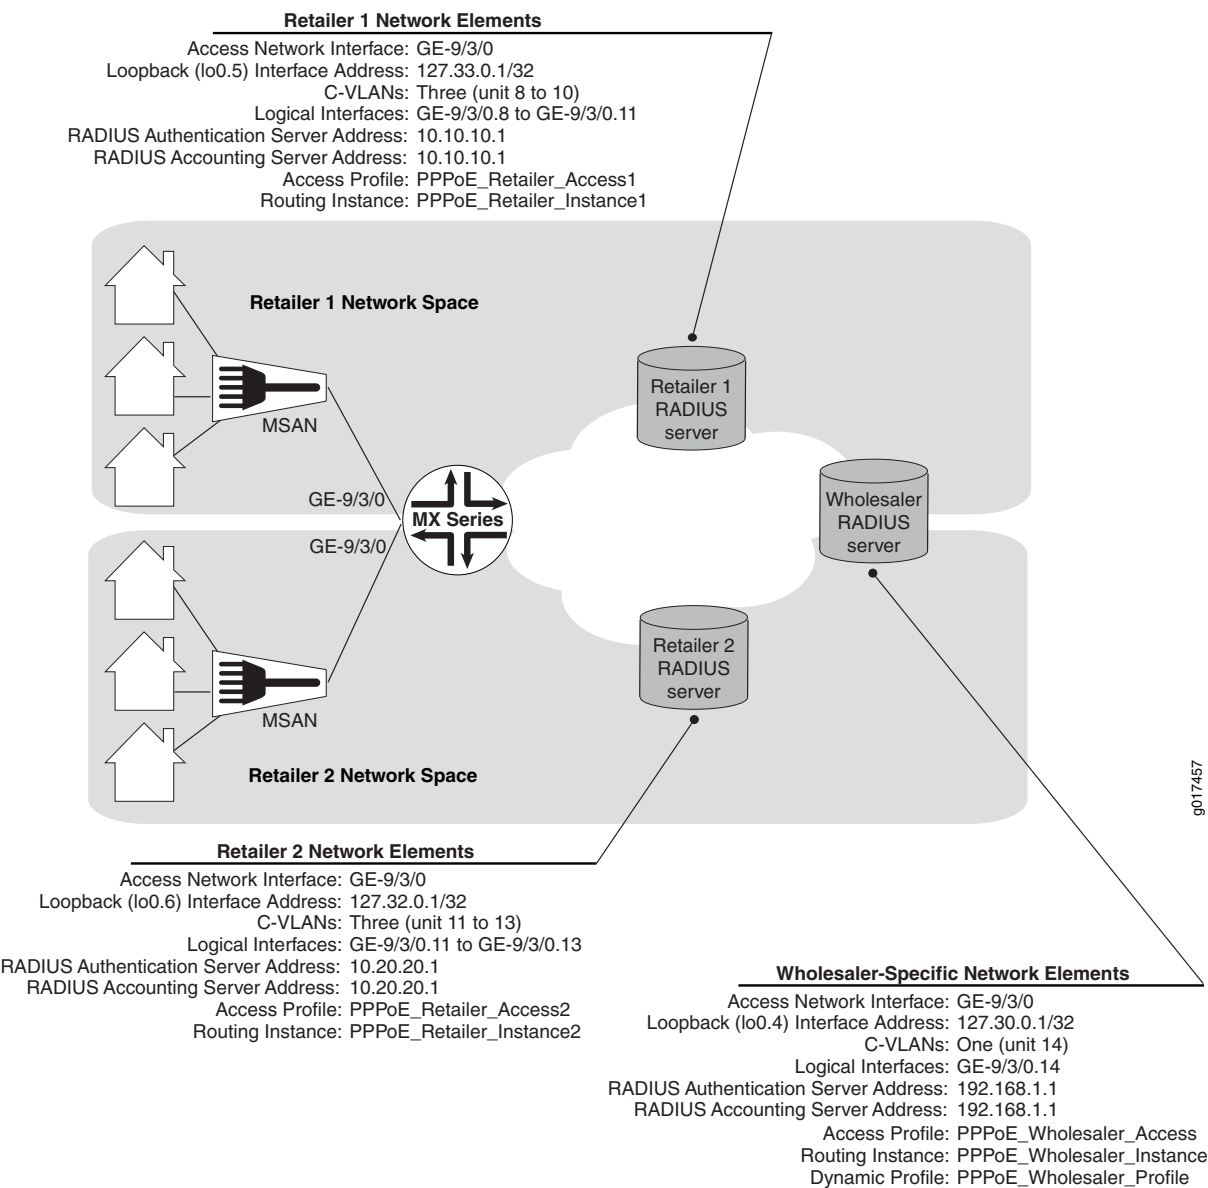
- Dynamic PPPoE interface creation.
- Static VLAN use only.
- AAA server assignment of subscribers to different routing instances within the same (default) logical system only.

RELATED DOCUMENTATION

PPPoE Layer 3 Wholesale Network Topology Overview

This configuration explains how to configure a simple PPPoE Layer 3 wholesale subscriber access network. This solution incorporates two retailers sharing resources on a wholesaler router. [Figure 6 on page 68](#) provides the reference topology for this configuration example.

Figure 6: PPPoE Layer 3 Wholesale Network Reference Topology



RELATED DOCUMENTATION

[Layer 2 and Layer 3 Wholesale Overview | 2](#)

[Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Topology and Configuration Elements | 8](#)

Configuring Loopback Interfaces for the PPPoE Layer 3 Wholesale Solution

You must configure loopback interfaces for use in the subscriber management access network. The loopback interfaces are 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 loopback interfaces:

1. Edit the loopback interface.

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

2. Edit the unit for the wholesale loopback interface.

```
[edit interfaces lo0]
user@host# edit unit 4
```

3. Edit the wholesale loopback interface family.

```
[edit interfaces lo0 unit 4]
user@host# edit family inet
```

4. Specify the wholesale loopback interface address.

```
[edit interfaces lo0 unit 4 family inet]
user@host# set address 127.30.0.1/32
```

5. (Optional) Specify the loopback interface address as the primary loopback interface.

```
[edit interfaces lo0 unit 4 family inet]
user@host# set address 127.30.0.2/32 primary
```

6. Edit the unit for a retail loopback interface.

```
[edit interfaces lo0]
user@host# edit unit 5
```

7. Edit the retail loopback interface family.

```
[edit interfaces lo0 unit 5]
user@host# edit family inet
```

8. Specify the retail loopback interface address.

```
[edit interfaces lo0 unit 5 family inet]
user@host# set address 127.33.0.1/32
```

9. (Optional) Specify the loopback interface address as the primary loopback interface.

```
[edit interfaces lo0 unit 5 family inet]
user@host# set address 127.33.0.2/32 primary
```

10. Repeat steps 7 through 10 for additional retailers, making sure to use unique unit and address values for each retailer loopback interface.

RELATED DOCUMENTATION

| *Junos OS Network Interfaces Library for Routing Devices*

Configuring Static Customer VLANs for the PPPoE Layer 3 Wholesale Network Solution

In this example configuration, the access interface (**ge-9/3/0**) 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 wholesaler and any access network subscribers.

To configure the customer VLANs:

1. Edit the access side interface.

```
[edit]
user@host# edit interfaces ge-9/3/0
```

2. Specify the use of flexible VLAN tagging.

```
[edit interfaces ge-9/3/0]
user@host# set flexible-vlan-tagging
```

3. Edit the interface unit for the wholesaler VLAN.

```
[edit interfaces ge-9/3/0]
user@host# edit unit 14
```

4. Specify the type of encapsulation that you want the wholesaler VLAN to use.

```
[edit interfaces ge-9/3/0 unit 14]
user@host# set encapsulation ppp-over-ether
```

5. (Optional) Specify that you want the wholesaler VLAN to use Proxy ARP.

```
[edit interfaces ge-9/3/0 unit 14]
user@host# set proxy-arp
```

6. Define a unique VLAN ID for the wholesaler VLAN.

```
[edit interfaces ge-9/3/0 unit 14]
user@host# set vlan-id 14
```

7. Specify the dynamic profile that you want the wholesaler VLAN to use.

```
[edit interfaces ge-9/3/0 unit 14]
user@host# set pppoe-underlying-options dynamic-profile PPPoE_Wholesaler_Profile
```

Configuring Access Components for the PPPoE Wholesale Network Solution

IN THIS SECTION

- [Configuring RADIUS Server Access | 72](#)
- [Configuring a PPPoE Wholesaler Access Profile | 73](#)
- [Configuring PPPoE Retailer Access Profiles | 74](#)

When configuring a wholesale network, you must configure several components globally. This configuration provides access to RADIUS servers (if used) that you want the wholesaler and any configured retailers to use globally. The access configuration includes the following general steps:

Configuring RADIUS Server Access

You can globally define any RADIUS servers in your network that either the wholesale access profile or retailer access profile can use. After you define the global RADIUS servers, you can specify specific RADIUS servers within individual access profiles.

To define RADIUS servers for profile access:

1. Access the **[edit access radius-server]** hierarchy level.

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

2. Specify the address and secret for any RADIUS servers in the network.

```
[edit access radius-server]
user@host# set 192.168.10.1 secret $ABC123$ABC123$ABC123
user@host# set 10.10.10.1 secret $ABC123$ABC123
```

SEE ALSO

RADIUS Servers and Parameters for Subscriber Access

Configuring a PPPoE Wholesaler Access Profile

You must define the network and interface over which you want subscribers to initially access the network with a wholesale access profile. When a subscriber attempts to access the network, the access profile provides initial access information including authentication and accounting values that the router uses for the accessing subscriber.

To define a wholesale access profile:

1. Create the wholesale access profile.

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

2. Specify the authentication methods for the profile and the order in which they are used.

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

3. Specify that you want to configure RADIUS support.

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

4. Specify the IP address of the RADIUS server used for authentication.

```
[edit access profile PPPoE_Wholesaler_Access radius]
user@host# set authentication-server 192.168.10.1
```

5. Specify the IP address of the RADIUS server used for accounting.

```
[edit access profile PPPoE_Wholesaler_Access radius]
user@host# set accounting-server 192.168.10.1
```

6. Configure any desired options for the RADIUS server.

See *RADIUS Servers and Parameters for Subscriber Access*.

7. Configure subscriber accounting (RADIUS accounting).

See *Configuring Per-Subscriber Session Accounting*.

SEE ALSO

[*Configuring Authentication and Accounting Parameters for Subscriber Access*](#)

[*Specifying the Authentication and Accounting Methods for Subscriber Access*](#)

[*RADIUS Servers and Parameters for Subscriber Access*](#)

[*Configuring Per-Subscriber Session Accounting*](#)

Configuring PPPoE Retailer Access Profiles

In this solution, subscribers are redirected to a networking space used by a specific retailer and defined by a unique routing instance. This method requires that you define the network and interface over which you want subscribers to access the network after being redirected by the wholesale access profile.

To define a retailer access profile:

1. Create the retailer access profile.

```
[edit]
user@host# edit access profile PPPoE_Retailer_Access1
```

2. Specify the authentication methods for the profile and the order in which they are used.

```
[edit access profile PPPoE_Retailer_Access1]
user@host# set authentication-order radius
```

3. Specify that you want to configure RADIUS support.

```
[edit access profile PPPoE_Retailer_Access1]
user@host# edit radius
```

4. Specify the IP address of the RADIUS server used for authentication.

```
[edit access profile PPPoE_Retailer_Access1 radius]
user@host# set authentication-server 10.10.10.1
```

5. Specify the IP address of the RADIUS server used for accounting.

```
[edit access profile PPPoE_Retailer_Access1 radius]
user@host# set accounting-server 10.10.10.1
```

6. Configure any desired options for the RADIUS server.

See *RADIUS Servers and Parameters for Subscriber Access*.

7. Configure subscriber accounting (RADIUS accounting).

See *Configuring Per-Subscriber Session Accounting*.

SEE ALSO

Configuring Authentication and Accounting Parameters for Subscriber Access

Specifying the Authentication and Accounting Methods for Subscriber Access

RADIUS Servers and Parameters for Subscriber Access

Configuring Per-Subscriber Session Accounting

Configuring Dynamic Profiles for the PPPoE Layer 3 Wholesale Network Solution

IN THIS SECTION

- [Configuring a Wholesale Dynamic Profile for use in the PPPoE Solution | 76](#)

A dynamic profile is a set of characteristics, defined in a type of template, that you can use to provide services for broadband applications. These services are assigned dynamically to interfaces as they access the network. When configuring dynamic profiles for the PPPoE Layer 3 wholesale network, you can choose to configure one dynamic profile to address all incoming subscribers or you can configure individual dynamic profiles for use by the different network management groups (that is, the wholesaler and any retailers). In fact, you can create multiple dynamic profiles that you can use to roll out different services and selectively apply those dynamic profiles to different subscriber groups as necessary.

In this solution example, one dynamic profile is created for use by the wholesaler when subscribers initially access the network. Subscribers are assigned by the wholesaler RADIUS server to a particular retailer routing instance and can then be redirected to that retailer network space.

Configuring a Wholesale Dynamic Profile for use in the PPPoE Solution

You can configure a basic access profile to initially manage PPPoE subscribers that access the network.

To configure a dynamic profile for use by the wholesaler:

1. Create a wholesale dynamic profile.

```
[edit]
user@host# edit dynamic-profiles PPPoE_Wholesaler_Profile
```

2. Define the dynamic routing instance variable in the dynamic profile.

```
[edit dynamic-profiles PPPoE_Wholesaler_Profile]
user@host# edit routing-instances $junos-routing-instance
```

3. Set the dynamic interface variable for the dynamic routing instance.

```
[edit dynamic-profiles PPPoE_Wholesaler_Profile routing-instances "$junos-routing-instance"]
user@host# set interface $junos-interface-name
```

4. Specify that you want to configure the **pp0** interface in the dynamic profile.

```
[edit dynamic-profiles PPPoE_Wholesaler_Profile]
user@host# edit interfaces pp0
```

5. Configure the unit for the **pp0** interface.

- a. Configure the variable for the unit number of the **pp0** interface.

The variable is dynamically replaced with the unit number that RADIUS supplies when the subscriber logs in.

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

- b. Configure PAP or CHAP (or both) to function on the interface.


```
[edit dynamic-profiles PPPoE_Wholesaler_Profile interfaces pp0 unit "$junos-interface-unit"]
user@host# set ppp-options chap pap
```

- c. Configure the variable for the underlying interface of the pp0 interfaces.

The variable is dynamically replaced with the underlying interface that RADIUS supplies when the subscriber logs in.

```
[edit dynamic-profiles PPPoE_Wholesaler_Profile interfaces pp0 unit "$junos-interface-unit"]
user@host# set pppoe-options underlying-interface $junos-underlying-interface
```

- d. Configure the router to act as a PPPoE server.

```
[edit dynamic-profiles PPPoE_Wholesaler_Profile interfaces pp0 unit "$junos-interface-unit"]
user@host# set pppoe-options server
```

6. (Optional) Modify the PPPoE keepalive interval.

```
[edit dynamic-profiles PPPoE_Wholesaler_Profile interfaces pp0 unit "$junos-interface-unit"]
user@host# set keepalives interval 15
```

7. Configure the family for the **pp0** interface.

- a. Specify that you want to configure the family.

NOTE: You can specify **inet** for IPv4 and **inet6** for IPv6. However, this solution provides the IPv4 configuration only.

```
[edit dynamic-profiles PPPoE_Wholesaler_Profile interfaces pp0 unit "$junos-interface-unit"]
user@host# edit family inet
```

- b. Configure the unnumbered address for the family.

```
[edit dynamic-profiles PPPoE_Wholesaler_Profile interfaces pp0 unit "$junos-interface-unit" family inet]
user@host# set unnumbered-address $junos-loopback-interface
```

Configuring Separate Routing Instances for PPPoE Service Retailers

As the owner of the system, the wholesaler uses the default routing instance. You must create separate routing instances for each individual retailer to keep routing information for individual retailers separate and to define any servers and forwarding options specific to each retailer.

To define a retailer routing instance:

1. Create the retailer routing instance.

```
[edit]
user@host# edit routing-instances PPPoE_Retailer_Instance1
```

2. Specify the routing instance type for the retailer.

```
[edit routing-instances "PPPoE_Retailer_Instance1"]
user@host# set instance-type vrf
```

3. Specify the access profile that you want the routing instance to use.

```
[edit routing-instances "PPPoE_Retailer_Instance1"]
user@host# set access-profile PPPoE_Retailer_Access1
```

4. Specify the interface that faces the Retailer1 RADIUS server.

```
[edit routing-instances "PPPoE_Retailer_Instance1"]
user@host# set interface ge-11/1/9.10
```

5. Specify the loopback interface unit for this routing instance.

```
[edit routing-instances "PPPoE_Retailer_Instance1"]
user@host# set interface lo0.5
```

NOTE: Loopback interfaces must be unique for each routing instance.

6. Specify an identifier to distinguish the VPN to which the route belongs.

```
[edit routing-instances "PPPoE_Retailer_Instance1"]
user@host# set route-distinguisher 1:1
```

7. Specify how routes are imported into the local PE router's VPN routing table from the remote PE router.

```
[edit routing-instances "PPPoE_Retailer_Instance1"]
user@host# set vrf-import policyImport
```

8. Specify which routes are exported from the local instance table to the remote PE router.

```
[edit routing-instances "PPPoE_Retailer_Instance1"]
user@host# set vrf-export policyExport
```

9. Repeat this procedure for other retailers.

RELATED DOCUMENTATION

Example: Wholesaler Dynamic Profile for a PPPoE Wholesale Network

This example specifies a dynamic profile name of *PPPoE_Wholesaler_Profile*, uses pp0 interfaces, and references the predefined input firewall filter.

```
PPPoE_Wholesaler_Profile {
  routing-instances {
    "$junos-routing-instance" {
      interface "$junos-interface-name";
    }
  }
  interfaces {
    pp0 {
      unit "$junos-interface-unit" {
        ppp-options {
          chap;
          pap;
        }
        pppoe-options {
```

```

        underlying-interface "$junos-underlying-interface";
        server;
    }
    keepalives interval 15;
    family inet {
        filter {
            input "$junos-input-filter";
            output "$junos-output-filter";
        }
        unnumbered-address "$junos-loopback-interface";
    }
}
}
}
}

```

RELATED DOCUMENTATION

[Configuring Dynamic Profiles for the PPPoE Layer 3 Wholesale Network Solution](#) | 75

Example: Retailer Routing Instances for a PPPoE Wholesale Network

```

routing-instances {
    PPPoE_Retailer_Instance1 {
        instance-type vrf;
        access-profile PPPoE_Retailer_Access1;
        interface ge-11/1/9.10;
        interface lo0.5;
        route-distinguisher 1:1;
        vrf-import policyImport;
        vrf-export policyExport;
    }
    Retailer_Instance2 {
        instance-type vrf;
        access-profile PPPoE_Retailer_Access2;
        interface ge-11/1/9.10;
        interface lo0.6;
        route-distinguisher 2:2;
        vrf-import policyImport;
    }
}

```

```
    vrf-export policyExport;  
  }  
}
```

RELATED DOCUMENTATION

| [Configuring Separate Routing Instances for PPPoE Service Retailers](#) | 78

3

PART

Configuring Layer 2 Wholesale Networks

[Subscriber Management Layer 2 Wholesale Overview](#) | 83

[Configuring Layer 2 Wholesale Networks](#) | 89

Subscriber Management Layer 2 Wholesale Overview

IN THIS CHAPTER

- [Layer 2 and Layer 3 Wholesale Overview | 83](#)
- [Wholesale Network Configuration Options and Considerations | 84](#)
- [RADIUS VSAs and Broadband Subscriber Management Wholesale Configuration Overview | 85](#)
- [Extensible Subscriber Services Manager | 87](#)

Layer 2 and Layer 3 Wholesale Overview

In general, wholesaling broadband services allows service providers to resell broadband services and allows other providers to deploy their own services over the incumbent network. There are different methods to partitioning an access network for resale. The two most common approaches are based on either Layer 2 or Layer 3 information. Wholesale access is the process by which the access network provider (the *wholesaler*) partitions the access network into separately manageable and accountable subscriber segments for resale to other network providers (or *retailers*).

In a Layer 3 wholesale configuration, you partition the wholesaler access network at the network layer or the subscriber IP component by associating the IP component with a distinct Layer 3 domain. In a Layer 2 wholesale configuration, you partition the access network at the subscriber circuit or customer VLAN (C-VLAN) by backhauling the connection through the service provider backbone network to the subscribing retailer network where the access traffic can be managed at higher layers.

In a Junos OS Dynamic Host Configuration Protocol (DHCP) or Point-to-Point Protocol over Ethernet (PPPoE) subscriber access configuration, wholesale partitioning is accomplished through the use of logical systems and routing instances within the router. Logical systems offer a stricter partitioning of routing resources than routing instances. The purpose behind the use of logical systems is to distinctly partition the physical router into separate administrative domains. This partitioning enables multiple providers to administer the router simultaneously, with each provider having access only to the portions of the configuration relevant to their logical system. Junos OS supports up to 15 named logical systems in addition to the default logical system (that is, **inet.0**). Unless otherwise specified in configuration, all interfaces belong to the default logical system.

NOTE: This Junos OS release supports the use of only the default logical system. Partitioning currently occurs through the use of separate routing instances.

A logical system can have one or more routing instances. Typically used in Layer 3 VPN scenarios, a routing instance does not have the same level of administrative separation as a logical system because it does not offer administrative isolation. However, the routing instance defines a distinct routing table, set of routing policies, and set of interfaces.

RELATED DOCUMENTATION

[Broadband Subscriber Management DHCPv4 Layer 3 Wholesale Topology and Configuration Elements | 8](#)

[Broadband Subscriber Management PPPoE Layer 3 Wholesale Topology and Configuration Elements | 66](#)

[Broadband Subscriber Management Layer 2 Wholesale Topology and Configuration Elements | 89](#)

Wholesale Network Configuration Options and Considerations

You can configure a wholesale network any number of ways using Juniper Networks hardware and Junos OS software. The general configuration options, and considerations for each, are provided in the following table:

Wholesale Configuration Options	Considerations
Fully Static (all interfaces, VLANs, and routing instances are configured statically)	Providing more control over retailer space and access, this option is more labor intensive and can require more detailed planning of the network, address allocation, and so on.
Static VLANs and Dynamic Demux Interfaces	Service VLANs are created statically and must be managed. Demux interfaces are dynamically created over the service VLANs. This option uses more logical interfaces; one for each VLAN and one for each dynamic demux interface that runs over each VLAN.

Wholesale Configuration Options	Considerations
Dynamic VLANs Only (dedicated customer VLANs for each subscriber)	<p>Dynamic (auto-sensed) VLANs are authenticated and installed in the correct non-default routing instance before DHCP is instantiated. This method helps to conserve logical interfaces by avoiding the need for additional logical interfaces being created for each demux interface.</p> <p>NOTE: In a customer VLAN model, each VLAN functions on a 1:1 basis for each customer (in this case, per household).</p>
Dynamic VLANs and Dynamic Demux Interfaces	<p>Allows for the greatest ease of use and flexibility in configuring subscribers, by enabling access over a service VLAN and targetting more service levels over individual, dynamically-created demux interfaces over the service VLAN. This option uses more logical interfaces; one for each VLAN and one for each demux interface that runs over each VLAN.</p>

RADIUS VSAs and Broadband Subscriber Management Wholesale Configuration Overview

You can use RADIUS to assign various values through the use of dynamic variables within dynamic profiles. However, the configuration of at least one of the two VSAs described in [Table 3 on page 6](#) is required for a wholesale network to function.

Table 5: Required Juniper Networks VSAs for the Broadband Subscriber Management Wholesale Network Solution

Attribute Number	Attribute Name	Description	Value
26-1	LSRI-Name	<p>Client logical system/routing instance membership name.</p> <p>Allowed only from RADIUS server for "default" logical system/routing instance membership.</p>	string: logical system:routing instance

Table 5: Required Juniper Networks VSAs for the Broadband Subscriber Management Wholesale Network Solution (continued)

Attribute Number	Attribute Name	Description	Value
26-25	Redirect-LSRI-Name	Client logical system/routing instance membership name indicating to which logical system/routing instance membership the request is redirected for user authentication.	string: logical system:routing instance

Specifying the **\$junos-routing-instance** dynamic variable in a dynamic profile triggers a RADIUS access-accept response of either the LSRI-Name VSA or the Redirect-LSRI-Name VSA. Returning an LSRI-Name attribute in the access-accept response provides the logical system and routing instance in which the logical interface is to be created and the router updates the session database with the specified routing instance value. Returning a Redirect-LSRI-Name attribute in the access-accept response results in the router immediately sending a second access-request message (sometimes referred to as a *double-dip*) to the RADIUS server specified by the logical system:routing instance attribute specified by the Redirect-LSRI-Name VSA.

NOTE: Attributes returned as a result of a second access-request message to the logical system/routing instance membership specified by the Redirect-LSRI-Name VSA override any prior attributes returned by initial access-accept responses to the default logical system/routing instance membership.

RELATED DOCUMENTATION

Juniper Networks VSAs Supported by the AAA Service Framework.

Extensible Subscriber Services Manager

IN THIS SECTION

- [Extensible Subscriber Services Manager Overview | 87](#)
- [Understanding the Dictionary File | 88](#)

Extensible Subscriber Services Manager Overview

Extensible Subscriber Services Manager (ESSM) is a background process that is part of the Intelligent Customer Extendable authentication, authorization, and accounting (ICE-AAA) framework, which supports customer extensible services for both business and residential subscribers. Services are classified as residential or business on the basis of the value specified for the RADIUS VSA (26-173)

ERX-Service-Activate-Type that is received in the Access-Accept message.

Extensible Subscriber Services Manager uses the ICE-AAA framework, which comprises a dictionary, operation scripts, and RADIUS vendor-specific attributes (VSAs), to create business services for subscribers without modifying Junos OS. Extensible Subscriber Services Manager supports only the **ERX-Activate** service type.

Using the Extensible Subscriber Services Manager, you can create business services using the following sources:

- The dictionary that refers to or invokes the operation scripts.
- The operation scripts that you use to create subscriber-specific configuration
- The VSAs that the RADIUS server sends that contain configuration values for provisioning services

SEE ALSO

[Understanding the Dictionary File | 88](#)

[show subscribers | 849](#)

[show subscribers summary | 896](#)

[Configuring Flat-File Accounting for Extensible Subscriber Services Management | 158](#)

Understanding the Dictionary File

The XML-based dictionary specifies the action to be taken by ESSMD when it receives a service request. The dictionary contains provisioning, deprovisioning, and operation scripts. ESSMD parses the dictionary file during initialization and stores the parsed information in the database. Extensible Subscriber Services Manager acts on the extensible-subscriber-service requests on the basis of the services configured in the dictionary file.

During a commit operation, essmd verifies the path and the filename of the dictionary file. If the path or the filename is invalid, the commit operation fails and the error is logged in a system log message. Restarting the daemon or performing a graceful Routing Engine switchover (GRES) operation forces essmd to use the new dictionary. Ensure that you always configure a valid dictionary for essmd.

When loading the dictionary file after a successful commit operation, essmd validates whether:

- There are errors in parsing the dictionary file.
- The operation scripts specified in the dictionary file are available on the router.
- Any active services are modified.

If the validation fails, an error is logged in a system log message, and essmd continues to use the existing version of the dictionary file. Use the **request services extensible-subscriber-services reload-dictionary** command to reload the dictionary file after resolving the errors.

SEE ALSO

<i>dictionary</i>
show extensible-subscriber-services dictionary 607
show extensible-subscriber-services dictionary attributes 612
show extensible-subscriber-services dictionary services 615

Configuring Layer 2 Wholesale Networks

IN THIS CHAPTER

- Broadband Subscriber Management Layer 2 Wholesale Topology and Configuration Elements | 89
- Layer 2 Wholesale Network Topology Overview | 91
- Configuring a Retail Dynamic Profile for Use in the Layer 2 Wholesale Solution | 92
- Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution | 94
- Configuring VLAN Interfaces for the Layer 2 Wholesale Solution | 97
- Configuring Encapsulation for Layer 2 Wholesale VLAN Interfaces | 98
- Configuring NNI ISP-Facing Interfaces for the Layer 2 Wholesale Solution | 100
- Configuring Direct ISP-Facing Interfaces for the Layer 2 Wholesale Solution | 101
- Configuring Separate Access Routing Instances for Layer 2 Wholesale Service Retailers | 102
- Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers | 105
- Configuring Access Components for the Layer 2 Wholesale Network Solution | 108
- Example: Retailer Dynamic Profile for a Layer 2 Wholesale Network | 110
- Example: Access Interface for a Layer 2 Wholesale Network | 111
- Example: Retailer Access Routing Instances for a Layer 2 Wholesale Network | 112
- Example: Retailer NNI ISP-Facing Interfaces for a Layer 2 Wholesale Network | 113
- Example: Retailer Direct ISP-Facing Interface for a Layer 2 Wholesale Network | 114

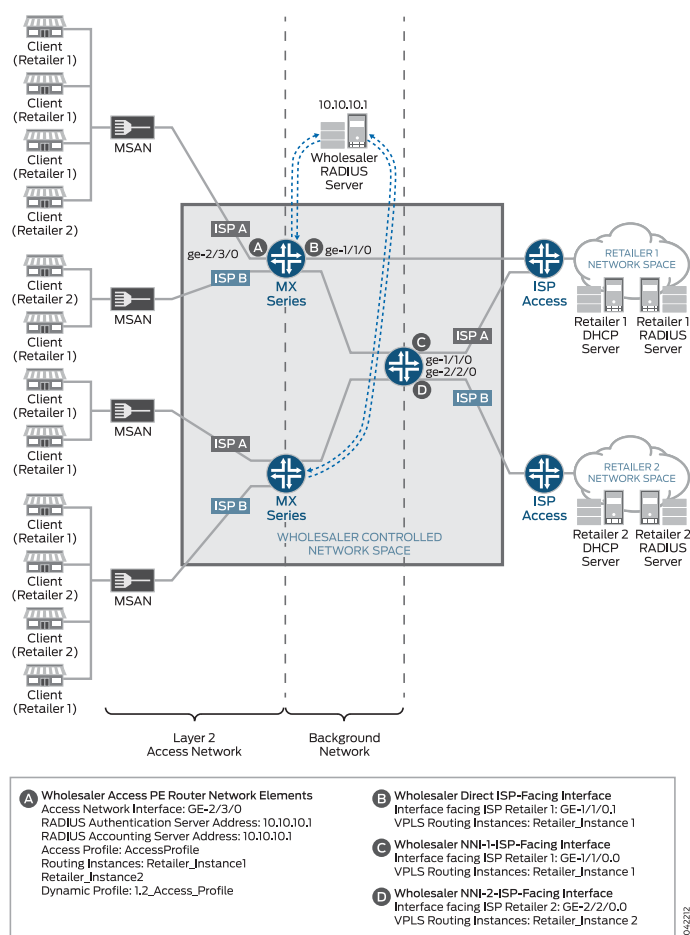
Broadband Subscriber Management Layer 2 Wholesale Topology and Configuration Elements

The network topology for the subscriber management Layer 2 wholesale solution includes configuring separate routing instances for individual retailers that use a portion of the router. This solution uses a Virtual Private LAN Service (VPLS) configuration.

Layer 2 wholesale networks are supported on MPC/MIC interfaces.

To explain the concept but limit complexity, this solution provides a configuration with one wholesaler and only two retailers. [Figure 7 on page 90](#) illustrates a basic Layer 2 wholesale topology model from which you can expand.

Figure 7: Basic Subscriber Management Layer 2 Wholesale Solution Topology



When you are configuring a Layer 2 wholesale network solution, the following configuration elements are required:

- Subscriber access dynamic VLAN configuration including dynamic profile configuration for retailer routing instances
- Routing instance configuration for individual retailers on provider edge (PE) routers and network-to-network interface (NNI) routers.
- VLAN interface configuration
- RADIUS server access configuration
- Core network configuration

RELATED DOCUMENTATION

[Layer 2 and Layer 3 Wholesale Overview | 2](#)

[Layer 2 Wholesale Network Topology Overview | 91](#)

Layer 2 Wholesale Network Topology Overview

This configuration explains how to configure a simple Layer 2 wholesale subscriber access network. This solution illustrates two Internet Service Provider (ISP) retailers sharing access to a wholesaler network. The wholesaler network contains a Layer 2 Network access router and two Virtual Private LAN Service (VPLS) network-to-network interface (NNI) routers.

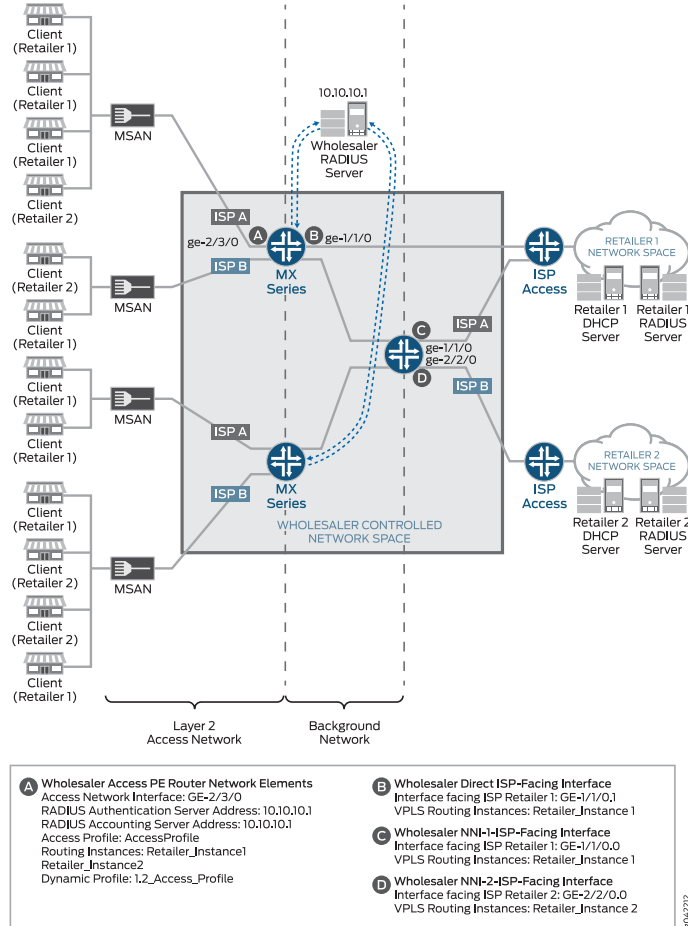
NOTE: You can have more than one ISP router connecting to a single VPLS NNI router with VPLS interfaces configured with routing instances specific to each different ISP-facing interfaces.

The example also shows two different connection options from one subscriber access router to one of the individual ISP access routers. One connection option uses an interface on the subscriber access router to connect directly to the ISP access router. Another connection option uses two routers: a subscriber access router and another NNI router that connects to the ISP access router.

NOTE: When using the NNI router connection option, use a standard BGP or MPLS configuration between the subscriber access routers and the edge router that connects to the ISP access routers. See the *BGP User Guide* for information about BGP configuration. See the *MPLS Applications User Guide* for information about MPLS configuration.

[Figure 8 on page 92](#) provides the reference topology for this configuration example.

Figure 8: Layer 2 Wholesale Network Reference Topology



RELATED DOCUMENTATION

[Layer 2 and Layer 3 Wholesale Overview | 2](#)

[Broadband Subscriber Management Layer 2 Wholesale Topology and Configuration Elements | 89](#)

Configuring a Retail Dynamic Profile for Use in the Layer 2 Wholesale Solution

To configure a dynamic profile for use with retailer access:

1. Create a retail dynamic profile.


```
[edit]
user@host# edit dynamic-profiles Subscriber_Profile_Retail1
```

2. Define the dynamic routing instance variable in the dynamic profile.

```
[edit dynamic-profiles Subscriber_Profile_Retail1]
user@host# edit routing-instances $junos-routing-instance
```

3. Set the dynamic interface variable for the dynamic routing instance.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 routing-instances "$junos-routing-instance"]
user@host# set interface $junos-interface-name
```

4. Define the dynamic interfaces variable for the dynamic profile.

```
[edit dynamic-profiles Subscriber_Profile_Retail1]
user@host# set interfaces $junos-interface-ifd-name
```

5. Define the dynamic interface unit variable for the dynamic profile.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces "$junos-interface-ifd-name"]
user@host# set unit $junos-interface-unit
```

6. (Optional) Define the VLAN encapsulation for the dynamic interfaces.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces "$junos-interface-ifd-name" unit
"$junos-interface-unit"]
user@host# set encapsulation vlan-vpls
```

NOTE: If you choose not to specify an encapsulation for the logical interface, you must specify encapsulation for the physical interface.

7. Define the VLAN tag variables for the dynamic profile:

NOTE: This solution example uses stacked VLAN tagging. However, you can also specify single-tag VLANs. For additional information about configuring dynamic VLANs, see the *Broadband Subscriber VLANs and Interfaces User Guide*.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces "$junos-interface-ifd-name" unit
"$junos-interface-unit"]
user@host# set vlan-tags outer $junos-stacked-vlan-id inner $junos-vlan-id
```

8. Define the input and output VLAN maps. See [“Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution” on page 94](#) for details.
9. Specify the unit family as **vpls** at the [edit dynamic-profiles *profile-name* interfaces “\$junos-interface-ifd-name” unit “\$junos-interface-unit” family] hierarchy level.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces "$junos-interface-ifd-name" unit
"$junos-interface-unit"]
user@host# set family vpls
```

Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution

Stacking and rewriting VLAN tags allows you to use an additional (outer) VLAN tag to differentiate between routers in the Layer 2 wholesale network. A frame can be received on an interface, or it can be internal to the system (as a result of the **input-vlan-map** statement).

You can configure rewrite operations to stack (**push**), remove (**pop**), or rewrite (**swap**) tags on single-tagged frames and dual-tagged frames. If a port is not tagged, rewrite operations are not supported on any logical interface on that port.

You can configure the following single-action VLAN rewrite operations:

- **pop**—Remove a VLAN tag from the top of the VLAN tag stack. The outer VLAN tag of the frame is removed.
- **push**—Add a new VLAN tag to the top of the VLAN stack. An outer VLAN tag is pushed in front of the existing VLAN tag.
- **swap**—Replace the inner VLAN tag of the incoming frame with a user-specified VLAN tag value.

You configure VLAN rewrite operations for logical interfaces in the input VLAN map for incoming frames and in the output VLAN map for outgoing frames.

You can include both the **input-vlan-map** and **output-vlan-map** statements at the **[edit dynamic-profiles profile-name interface “\$junos-interface-ifd-name” unit ” \$junos-interface-unit]** hierarchy level.

The type of VLAN rewrite operation permitted depends upon whether the frame is single-tagged or dual-tagged. [Table 6 on page 95](#) shows supported rewrite operations and whether they can be applied to single-tagged frames or dual-tagged frames. The table also indicates the number of tags being added or removed during the operation.

Table 6: Rewrite Operations on Single-Tagged and Dual-Tagged Frames

Rewrite Operation	Single-Tagged	Dual-Tagged	Number of Tags
pop	Yes	Yes	- 1
push	Yes	Yes	+1
swap	Yes	Yes	0

Depending on the VLAN rewrite operation, you configure the rewrite operation for the interface in the input VLAN map, the output VLAN map, or both. [Table 7 on page 95](#) shows what rewrite operation combinations you can configure. “None” means that no rewrite operation is specified for the VLAN map.

Table 7: Applying Rewrite Operations to VLAN Maps

Input VLAN Map	Output VLAN Map			
	none	push	pop	swap
none	Yes	No	No	Yes
push	No	No	Yes	No
pop	No	Yes	No	No
swap	Yes	No	No	Yes

To configure the input VLAN map:

NOTE: You configure the **input-vlan-map** statement only when there is a need either to push an outer tag on a single-tagged subscriber packet or to modify the outer tag in a subscriber dual-tagged packet.

1. Include the **input-vlan-map** statement.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces "$junos-interface-ifd-name" unit
 "$junos-interface-unit"]
user@host# edit input-vlan-map
```

2. Specify the action that you want the input VLAN map to take.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces "$junos-interface-ifd-name" unit
 "$junos-interface-unit" input-vlan-map]
user@host# set push
```

3. Include the **vlan-id** statement along with the **\$junos-vlan-map-id** dynamic variable.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces "$junos-interface-ifd-name" unit
 "$junos-interface-unit" input-vlan-map]
user@host# set vlan-id $junos-vlan-map-id
```

To configure the output VLAN map:

NOTE: You configure the **output-vlan-map** statement only when there is a need to either pop or modify the outer tag found in a dual-tagged packet meant for the subscriber.

1. Include the **output-vlan-map** statement.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces "$junos-interface-ifd-name" unit
 "$junos-interface-unit"]
user@host# edit output-vlan-map
```

2. Specify the action that you want the output VLAN map to take.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces "$junos-interface-ifd-name" unit
"$junos-interface-unit" output-vlan-map]
user@host# set pop
```

You must know whether the VLAN rewrite operation is valid and is applied to the input VLAN map or the output VLAN map. You must also know whether the rewrite operation requires you to include statements to configure the inner and outer tag protocol identifiers (TPIDs) and inner and outer VLAN IDs in the input VLAN map or output VLAN map. For information about configuring inner and outer TPIDs and inner and outer VLAN IDs, see *Configuring Inner and Outer TPIDs and VLAN IDs*.

Configuring VLAN Interfaces for the Layer 2 Wholesale Solution

Clients access the Layer 2 Wholesale network through a specific interface. After they access this interface, and when they are authenticated, VLANs are dynamically created to carry the client traffic.

To configure a VLAN interface for dynamic access of clients:

1. Access the physical interface that you want to use for dynamically creating VLAN interfaces.

```
[edit interfaces]
user@host# edit interfaces ge-2/3/0
```

2. Specify the desired VLAN tagging.

NOTE: This example uses flexible VLAN tagging to simultaneously support transmission of 802.1Q VLAN single-tag and dual-tag frames on logical interfaces on the same Ethernet port.

```
[edit interfaces ge-2/3/0]
user@host# set flexible-vlan-tagging
```

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

```
[edit interfaces ge-2/3/0]
user@host# edit auto-configure
```

4. Specify that you want to configure single VLANs.

```
[edit interfaces ge-2/3/0 auto-configure]
user@host# edit vlan-ranges
```

5. Define the VLAN ranges for the configuration.

```
[edit interfaces ge-2/3/0 auto-configure vlan-ranges]
user@host# set ranges any, any
```

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

```
[edit interfaces ge-2/3/0 auto-configure vlan-ranges]
user@host# set dynamic-profile Subscriber_Profile_Retail1
```

7. Specify that any type of VLAN Ethernet packet is accepted by the interface.

```
[edit interfaces ge-2/3/0 auto-configure vlan-ranges dynamic-profile "Subscriber_Profile_Retail1"]
user@host# set accept any
```

8. Repeat steps for any other interfaces that you want to use for creating VLANs.

9. Specify the encapsulation type for the VLAN interfaces.

```
[edit interfaces ge-2/3/0]
user@host# edit encapsulation flexible-ethernet-services
```

RELATED DOCUMENTATION

| [Configuring Encapsulation for Layer 2 Wholesale VLAN Interfaces](#) | 98

Configuring Encapsulation for Layer 2 Wholesale VLAN Interfaces

Each dynamic VLAN interface in a Layer 2 wholesale network must use encapsulation. You can configure encapsulation dynamically for each VLAN interface by using the **encapsulation** statement at the **[edit dynamic-profiles *profile-name* interface "\$junos-interface-ifd-name" unit "\$junos-interface-unit"]** hierarchy level or configure encapsulation for the physical interfaces at the **[edit interfaces *interface-name*]** hierarchy

level for each dynamically created VLAN interface to use. However, how you choose to configure (or not configure) encapsulation at the [edit dynamic-profiles *profile-name* interface “\$junos-interface-ifd-name” unit “\$junos-interface-unit”] hierarchy level affects how you configure encapsulation at the [edit interfaces *interface-name*] hierarchy level.

Table 8 on page 99 provides the valid encapsulation combinations for both dynamic profiles and physical interfaces in the Layer 2 wholesale network.

Table 8: Encapsulation Combinations for Layer 2 Wholesale Interfaces

Dynamic Profile Encapsulation	Physical Interface Encapsulation	Usage Notes
vlan-vpls	vlan-vpls	Using the vlan-vpls encapsulation type in both the dynamic profile and when configuring the physical interface limits the VLAN ID value to a number greater than or equal to 512.
vlan-vpls	flexible-ethernet-services	Using the flexible-ethernet-services encapsulation type removes any VLAN ID value limitation.
vlan-vpls	extended-vlan-vpls	The extended-vlan-vpls encapsulation type can support multiple TPIDs. Using this encapsulation type removes any VLAN ID value limitation.
No encapsulation type	extended-vlan-vpls	The extended-vlan-vpls encapsulation type can support multiple TPIDs. Using this encapsulation type removes any VLAN ID value limitation.

To configure encapsulation for Layer 2 wholesale VLAN interfaces:

1. (Optional) Define the VLAN encapsulation for the dynamic interfaces.

```
[edit dynamic-profiles Subscriber_Profile_Retail1 interfaces “$junos-interface-ifd-name” unit
“$junos-interface-unit”]
user@host# set encapsulation encapsulation-type
```

2. Specify the encapsulation type for the physical VLAN interface.

```
[edit interfaces ge-2/3/0]
user@host# edit encapsulation encapsulation-type
```

NOTE: If you choose not to specify an encapsulation for the logical interface, you must specify **extended-vlan-vpls** encapsulation for the physical interface.

RELATED DOCUMENTATION

[Configuring a Retail Dynamic Profile for Use in the Layer 2 Wholesale Solution | 92](#)

[Configuring VLAN Interfaces for the Layer 2 Wholesale Solution | 97](#)

Configuring NNI ISP-Facing Interfaces for the Layer 2 Wholesale Solution

You must configure separate, ISP-facing interfaces on each NNI ISP-facing router that connect to individual retailer ISP access routers in the Layer 2 Wholesale solution.

NOTE: On the network-to-network (NNI) or egress interfaces of provider edge (PE) routers, you cannot configure the **inner-range vid1—vid2** option with the **vlan-tags** statement for ISP-facing interfaces.

To configure an NNI ISP-facing interface:

1. Access the physical interface that you want to use to access the retailer ISP network.

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

2. Specify the encapsulation type for the VLAN interfaces.

```
[edit interfaces ge-1/1/0]
user@host# edit encapsulation ethernet-vpls
```

3. Specify the interface unit that you want ISP clients to use.

```
[edit interfaces ge-1/1/0]
```



```
user@host# edit unit 0
```

4. Repeat these steps for any other NNI ISP-facing interfaces that you want to use. In this example, you must also configure interface **ge-2/2/0.0**.

RELATED DOCUMENTATION

[Configuring Direct ISP-Facing Interfaces for the Layer 2 Wholesale Solution | 101](#)

[Configuring Separate Access Routing Instances for Layer 2 Wholesale Service Retailers | 102](#)

Configuring Direct ISP-Facing Interfaces for the Layer 2 Wholesale Solution

When connecting a subscriber access router directly to an ISP access router, you must define any ISP-facing interfaces that connect to the retailer ISP access routers as core-facing interfaces.

To configure a direct ISP-facing interface:

1. Access the physical interface that you want to use to access the retailer ISP network.

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

2. Specify the encapsulation type for the VLAN interfaces.

```
[edit interfaces ge-1/1/0]
user@host# edit encapsulation ethernet-vpls
```

3. Specify the interface unit that you want ISP clients to use.

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

4. Specify the unit family.

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

5. Define the interface as core-facing to ensure that the network does not improperly treat the interface as a client interface..

```
[edit interfaces ge-1/1/0 unit 1 family vpls]
user@host# set core-facing
```

6. Repeat steps for any other direct ISP-facing interfaces that you want to use..

RELATED DOCUMENTATION

[Configuring NNI ISP-Facing Interfaces for the Layer 2 Wholesale Solution | 100](#)

[Configuring Separate Access Routing Instances for Layer 2 Wholesale Service Retailers | 102](#)

Configuring Separate Access Routing Instances for Layer 2 Wholesale Service Retailers

As the owner of the system, the wholesaler uses the default routing instance. You must create separate routing instances for each individual retailer to keep routing information for individual retailers separate and to define any servers and forwarding options specific to each retailer.

When creating separate routing instances, it is important to understand the role that the router plays in the Layer 2 Wholesale network and specify that role (either access or NNI) in the routing instance configuration. If the router connects directly to an ISP network (or ISP-controlled device), you must configure the routing instances as an NNI routing instance. See “[Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers](#)” on page 105.

To define an access retailer routing instance:

1. Create the retailer routing instance.

```
[edit]
user@host# edit routing-instances RetailerInstance1
```

2. Specify the VLAN model that you want the retailer to follow.

```
[edit routing-instances RetailerInstance1]
user@host# set vlan-model one-to-one
```

3. Specify the role that you want the routing instance to take.

```
[edit routing-instances RetailerInstance1]  
user@host# set instance-role access
```

4. Specify the routing instance type for the retailer.

```
[edit routing-instances RetailerInstance1]  
user@host# set instance-type l2backhaul-vpn
```

5. Specify the access interface for the retailer.

```
[edit routing-instances RetailerInstance1]  
user@host# set interface ge-2/3/0.0
```

6. Specify that access ports in this VLAN domain do not forward packets to each other.

```
[edit routing-instances RetailerInstance1]  
user@host# set no-local-switching
```

7. Specify a unique identifier attached to a route that enables you to distinguish to which VPN the route belongs.

```
[edit routing-instances RetailerInstance1]  
user@host# set route-distinguisher 10.10.1.1:1
```

8. (Optional) Specify a VRF target community.

```
[edit routing-instances RetailerInstance1]  
user@host# set vrf-target target:100:1
```

NOTE: The purpose of the **vrf-target** statement is to simplify the configuration by allowing you to configure most statements at the **[edit routing-instances]** hierarchy level.

9. Define the VPLS protocol for the routing instance.

a. Access the routing instance **protocols** hierarchy.

```
[edit routing-instances RetailerInstance1]  
user@host# edit protocols
```

b. Enable VPLS on the routing instance.

```
[edit routing-instances RetailerInstance1 protocols]  
user@host# edit vpls
```

c. Specify the maximum number of sites allowed for the VPLS domain.

```
[edit routing-instances RetailerInstance1 protocols vpls]  
user@host# set site-range 10
```

d. Specify the size of the VPLS MAC address table for the routing instance.

```
[edit routing-instances RetailerInstance1 protocols vpls]  
user@host# set mac-table-size 6000
```

e. Specify the maximum number of MAC addresses that can be learned by the VPLS routing instance.

```
[edit routing-instances RetailerInstance1 protocols vpls]  
user@host# set interface-mac-limit 2000
```

f. (Optional) Specify the **no-tunnel-services** statement if the router does not have a Tunnel Services PIC.

```
[edit routing-instances RetailerInstance1 protocols vpls]  
user@host# set no-tunnel-services
```

g. Specify a site name.

```
[edit routing-instances RetailerInstance1 protocols vpls]  
user@host# set site A-PE
```

- h. Specify a site identifier.

```
[edit routing-instances RetailerInstance1 protocols vpls site A-PE]
user@host# set site-identifier 1
```

10. Repeat this procedure for other retailers. In this example, you must configure a routing instance for Retailer 2.

RELATED DOCUMENTATION

Configuring the VPLS Site Name and Site Identifier

[Configuring NNI ISP-Facing Interfaces for the Layer 2 Wholesale Solution | 100](#)

[Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers | 105](#)

Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers

As the owner of the system, the wholesaler uses the default routing instance. You must create separate routing instances for each individual retailer to keep routing information for individual retailers separate and to define any servers and forwarding options specific to each retailer.

When creating separate routing instances, it is important to understand the role that the router plays in the Layer 2 Wholesale network and specify that role (either access or NNI) in the routing instance configuration. If the router connects to the access portion of the network (for example, to an MSAN device), you must configure the routing instances as an access routing instance. See [“Configuring Separate Access Routing Instances for Layer 2 Wholesale Service Retailers” on page 102](#).

To define a retailer routing instance:

1. Create the retailer routing instance.

```
[edit]
user@host# edit routing-instances RetailerInstance1
```

2. Specify the VLAN model that you want the retailer to follow.

```
[edit routing-instances RetailerInstance1]
```

```
user@host# set vlan-model one-to-one
```

3. Specify the role that you want the routing instance to take.

```
[edit routing-instances RetailerInstance1]
user@host# set instance-role nni
```

4. Specify the routing instance type for the retailer.

```
[edit routing-instances RetailerInstance1]
user@host# set instance-type l2backhaul-vpn
```

5. Define the NNI ISP-facing interface for this retailer.

```
[edit routing-instances RetailerInstance1]
user@host# set interface ge-1/1/0.0
```

6. Specify that access ports in this VLAN domain do not forward packets to each other.

```
[edit routing-instances RetailerInstance1]
user@host# set no-local-switching
```

7. Specify a unique identifier attached to a route that enables you to distinguish to which VPN the route belongs.

```
[edit routing-instances RetailerInstance1]
user@host# set route-distinguisher 10.10.10.1:1
```

8. (Optional) Specify a VRF target community.

```
[edit routing-instances RetailerInstance1]
user@host# set vrf-target target:100:1
```

NOTE: The purpose of the **vrf-target** statement is to simplify the configuration by allowing you to configure most statements at the **[edit routing-instances]** hierarchy level.

9. Define the VPLS protocol for the routing instance.

a. Access the routing instance **protocols** hierarchy.

```
[edit routing-instances RetailerInstance1]
user@host# edit protocols
```

b. Enable VPLS on the routing instance.

```
[edit routing-instances RetailerInstance1 protocols]
user@host# edit vpls
```

c. Specify the maximum number of sites allowed for the VPLS domain.

```
[edit routing-instances RetailerInstance1 protocols vpls]
user@host# set site-range 1000
```

d. (Optional) Specify the **no-tunnel-services** statement if the router does not have a Tunnel Services PIC.

```
[edit routing-instances RetailerInstance1 protocols vpls]
user@host# set no-tunnel-services
```

e. Specify a site name.

```
[edit routing-instances RetailerInstance1 protocols vpls]
user@host# set site A-PE
```

f. Specify a site identifier.

```
[edit routing-instances RetailerInstance1 protocols vpls site A-PE]
user@host# set site-identifier 1
```

g. Define the connectivity of the VPLS routing instance as **permanent** to keep the VPLS connection up until specifically taken down.

```
[edit routing-instances RetailerInstance1 protocols vpls]
user@host# set connectivity-type permanent
```

10. Repeat this procedure for other retailers.

RELATED DOCUMENTATION

[Configuring the VPLS Site Name and Site Identifier](#)

[Configuring VLAN Interfaces for the Layer 2 Wholesale Solution | 97](#)

[Configuring Separate Access Routing Instances for Layer 2 Wholesale Service Retailers | 102](#)

Configuring Access Components for the Layer 2 Wholesale Network Solution

IN THIS SECTION

- [Configuring RADIUS Server Access | 108](#)
- [Configuring a Layer 2 Wholesaler Access Profile | 109](#)

When configuring a wholesale network, you must configure several components globally. This configuration provides access to RADIUS servers (if used) that you want the wholesaler and any configured retailers to use globally. The access configuration includes the following general steps:

Configuring RADIUS Server Access

You can globally define any RADIUS servers in your network that either the wholesale access profile or retailer access profile can use. After you define the global RADIUS servers, you can specify specific RADIUS servers within individual access profiles.

To define RADIUS servers for profile access:

1. Access the **[edit access radius-server]** hierarchy level.

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

2. Specify the address and secret for any RADIUS servers in the network.


```
[edit access radius-server]
user@host# set 192.168.10.1 secret $ABC123$ABC123$ABC123
user@host# set 10.10.10.1 secret $ABC123$ABC123
```

SEE ALSO

RADIUS Servers and Parameters for Subscriber Access

Configuring a Layer 2 Wholesaler Access Profile

You must define the network and interface over which you want subscribers to initially access the network with a wholesale access profile. When a subscriber attempts to access the network, the access profile provides initial access information including authentication and accounting values that the router uses for the accessing subscriber.

To define a wholesale access profile:

1. Create the wholesale access profile.

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

2. Specify the authentication methods for the profile and the order in which they are used.

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

3. Specify that you want to configure RADIUS support.

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

4. Specify the IP address of the RADIUS server used for authentication.

```
[edit access profile AccessProfile radius]
user@host# set authentication-server 10.10.10.1
```

5. Specify the IP address of the RADIUS server used for accounting.

```
[edit access profile AccessProfile radius]
user@host# set accounting-server 10.10.10.1
```

6. Configure any desired options for the RADIUS server.

See *Configuring Access Profile Options for Interactions with RADIUS Servers*.

7. Configure subscriber accounting (RADIUS accounting).

See *Configuring Per-Subscriber Session Accounting*.

SEE ALSO

Configuring Authentication and Accounting Parameters for Subscriber Access

Specifying the Authentication and Accounting Methods for Subscriber Access

RADIUS Servers and Parameters for Subscriber Access

Configuring Per-Subscriber Session Accounting

Example: Retailer Dynamic Profile for a Layer 2 Wholesale Network

```
dynamic-profiles {
  Subscriber_Profile_Retail1 {
    routing-instances {
      "$junos-routing-instance" {
        interface "$junos-interface-name";
      }
    }
    interfaces {
      "$junos-interface-ifd-name" {
        unit "$junos-interface-unit" {
          encapsulation vlan-vpls;
          vlan-tags outer "$junos-stacked-vlan-id" inner "$junos-vlan-id";
          input-vlan-map {
            swap;
            vlan-id "$junos-vlan-map-id";
          }
          output-vlan-map swap;
          family vpls;
        }
      }
    }
  }
}
```

```

    }
  }
}

```

RELATED DOCUMENTATION

[Layer 2 and Layer 3 Wholesale Overview | 2](#)

[Layer 2 Wholesale Network Topology Overview | 91](#)

[Configuring a Retail Dynamic Profile for Use in the Layer 2 Wholesale Solution | 92](#)

Example: Access Interface for a Layer 2 Wholesale Network

```

interfaces {
  ge-2/3/0 {
    flexible-vlan-tagging;
    auto-configure {
      stacked-vlan-ranges {
        dynamic-profile Subscriber_Profile_Retail1 {
          accept any;
          ranges {
            any,any;
          }
        }
      }
      access-profile AccessProfile;
    }
  }
  encapsulation flexible-ethernet-services;
}

```

RELATED DOCUMENTATION

[Layer 2 and Layer 3 Wholesale Overview | 2](#)

[Layer 2 Wholesale Network Topology Overview | 91](#)

[Configuring VLAN Interfaces for the Layer 2 Wholesale Solution | 97](#)

Example: Retailer Access Routing Instances for a Layer 2 Wholesale Network

You need to create a routing instance for each retailer to keep routing information for different retailers separate and to define servers and forwarding options specific to each retailer.

There are two types of routing instances that you can create: access or NNI. The following code snippets show how to configure separate access routing instances for two retailers: `Retailer_Instance1` and `Retailer_Instance2`.

```
routing-instances {
  Retailer_Instance1 {
    vlan-model one-to-one;
    instance-role access;
    instance-type l2backhaul-vpn;
    interface ge-1/1/0.0
    no-local-switching;
    route-distinguisher 10.10.1.1:1;
    vrf-target target:100:1;
    protocols {
      vpls {
        site-range 10;
        mac-table-size {
          6000;
        }
        interface-mac-limit {
          2000;
        }
        no-tunnel-services;
        site A-PE {
          site-identifier 1;
        }
      }
    }
  }
  Retailer_Instance2 {
    vlan-model one-to-one;
    instance-role access;
    instance-type l2backhaul-vpn;
    interface ge-2/2/0.0
    no-local-switching;
    route-distinguisher 10.10.1.1:2;
    vrf-target target:300:1;
    protocols {
      vpls {
```

```

        site-range 1000;
        no-tunnel-services;
        site A-PE {
            site-identifier 1;
        }
    }
}
}
}

```

RELATED DOCUMENTATION

[Layer 2 and Layer 3 Wholesale Overview | 2](#)

[Layer 2 Wholesale Network Topology Overview | 91](#)

[Configuring Separate Access Routing Instances for Layer 2 Wholesale Service Retailers | 102](#)

Example: Retailer NNI ISP-Facing Interfaces for a Layer 2 Wholesale Network

```

interfaces {
    ge-1/1/0 {
        description Retailer 1 NNI ISP-facing interface;
        encapsulation ethernet-vpls;
        unit 0 {
        }
    }
    interfaces {
        ge-2/2/0 {
            description Retailer 2 NNI ISP-facing interface;
            encapsulation ethernet-vpls;
            unit 0;
        }
    }
}

```

RELATED DOCUMENTATION

[Layer 2 and Layer 3 Wholesale Overview | 2](#)

[Layer 2 Wholesale Network Topology Overview | 91](#)

Example: Retailer Direct ISP-Facing Interface for a Layer 2 Wholesale Network

```
interfaces {  
  ge-1/1/0 {  
    description Retailer 1 Direct ISP-facing interface;  
    encapsulation ethernet-vpls;  
    unit 1  
      family vpls {  
        core-facing;  
      }  
    }  
  }  
}
```

RELATED DOCUMENTATION

[Layer 2 and Layer 3 Wholesale Overview | 2](#)

[Layer 2 Wholesale Network Topology Overview | 91](#)

[Configuring Direct ISP-Facing Interfaces for the Layer 2 Wholesale Solution | 101](#)

4

PART

Configuring ANCP-Triggered Layer 2 Wholesale Services

[ANCP-Triggered Layer 2 Wholesale Service Overview | 116](#)

[Configuring ANCP-Triggered Layer 2 Wholesale Services | 136](#)

[Configuring Flat-File Accounting for Layer 2 Wholesale Services | 149](#)

[Configuring Five-Level and Four-Level Heterogeneous Networks | 168](#)

ANCP-Triggered Layer 2 Wholesale Service Overview

IN THIS CHAPTER

- [Layer 2 Wholesale with ANCP-Triggered VLANs Overview | 116](#)

Layer 2 Wholesale with ANCP-Triggered VLANs Overview

IN THIS SECTION

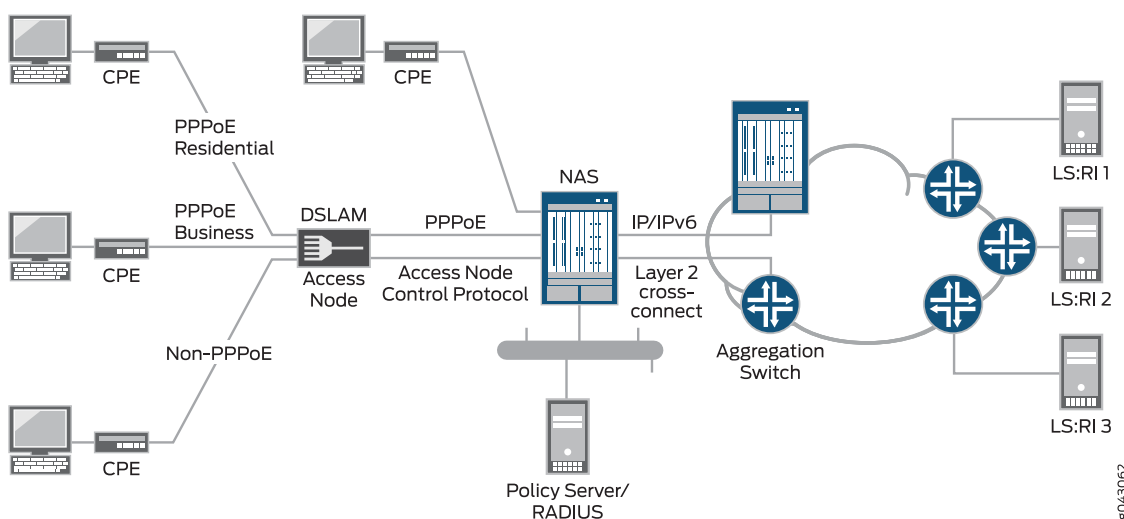
- [RADIUS Authorization for ANCP-Triggered VLANs | 119](#)
- [Instantiation of an ANCP-Triggered, Autosensed, Dynamic VLAN | 119](#)
- [Weighted Load Balancing for Subscriber Sessions over Eligible Core-Facing Physical Interfaces | 120](#)
- [RADIUS Interim Accounting Updates | 121](#)
- [Removal of the Layer 2 Wholesale Service | 122](#)
- [Interactions Between In-Band and Out-of-Band VLAN Autosensing | 123](#)
- [Migration of Subscriber Ownership from Wholesaler to Retailer | 125](#)
- [Migration of Subscriber Ownership from Retailer to Wholesaler | 126](#)
- [Migration of Subscriber Ownership Between Retailers | 127](#)
- [Modification of the Access Line Identifier or Port VLAN Identifier | 127](#)
- [Disconnecting PPPoE Sessions and Automatically Attempting Reconnection as Layer 2 Wholesale Sessions | 128](#)
- [Consequences of a State Transition in the Access-Facing Physical Interface | 129](#)
- [Consequences of a State Transition from Up to Down in the Core-Facing Physical Interface | 131](#)
- [Consequences of a State Transition from Down to Up in the Core-Facing Physical Interface | 132](#)
- [Loss of ANCP TCP Adjacency | 133](#)

The conventional mechanism for triggering autosensed dynamic VLANs relies on access line attributes provided by PPPoE or DHCP traffic in upstream control packets. Packets of a specified type are exceptioned

and authorization depends on fields extracted from the packet as specified in a dynamic profile assigned to the autosensed VLAN range. However, for some wholesale networks, the traffic might not be PPPoE or DHCP. In this case, a different mechanism is required.

Figure 9 on page 117 shows a sample topology with direct connections between the wholesaler's BNG and the NSP (network service provider) routers for the retailers. Each retailer's network resides in a dedicated routing instance. The wholesaler uses Layer 2 cross-connects to implement the retail networks with 1:1 autosensed, dynamic VLANs and VLAN tag swapping. Core-facing physical interfaces are dedicated to forwarding subscriber connections to the retailer's router. The traffic for an entire outer VLAN can be wholesaled this way. This direct-connect model supports any combination of wholesaler-owned and wholesaled connections for the entire access-facing VLAN range.

Figure 9: Sample Layer 2 Wholesale Access Topology



A wholesaler providing Layer 2 bitstream access to NSP partners might use this model. Bitstream access enables retailers to offer bidirectional transmission of broadband data and other high-speed services directly to customers across the wholesaler's network. In this topology, the PPPoE residential and subscriber customers are retained by the wholesaler (access provider). The non-PPPoE connections (here multiple connections and subscribers are represented by a single line) can be wholesaled to retail NSPs.

In this model, dynamic VLAN detection and creation for the wholesaled connections do not use in-band control packets. Instead, they rely on an out-of-band protocol, ANCP. ANCP Port Up messages both announce to the ANCP agent on the BNG that new access lines are operational and provide updates about previously announced lines. The messages include ANCP DSL attributes that correspond to Juniper Networks DSL VSAs and DSL Forum VSAs.

Starting in Junos OS Release 16.1R4, you can configure the ANCP agent to trigger the creation of an autosensed VLAN when the ANCP agent receives a Port Up message where the DSL-Line-State attribute has a value of Showtime. The Showtime state indicates that ports are configured, the subscriber is connected,

and the DSL modem is online and ready to transfer data. The other possible values of the attribute, Idle and Silent, are ignored for this purpose and are used by the ANCP agent only to update the ANCP session database (SDB).

During VLAN authorization, RADIUS determines which traffic belongs to the access provider's own subscribers and which belongs to the wholesale customer (retail NSP) based on identification of the subscriber's access line by the agent remote identifier.

When the ANCP agent receives the Port Up message, the agent triggers the autoconfiguration daemon, `autoconfd`, to initiate the VLAN detection, authorization, and creation processes. Those processes require the following information:

- Three ANCP subscriber access loop attributes (TLVs) that identify the access line and are conveyed in the Port Up message:
 - Access-Loop-Circuit-ID—Access loop circuit identifier used by the ANCP agent to determine which logical interface or interface set corresponds to the subscriber; corresponds to the Juniper Networks Acc-Loop-Cir-ID VSA (26-110).
 - Access-Loop-Remote-ID—Unique identifier of the access line; corresponds to the Juniper Networks Acc-Loop-Remote-ID VSA (26-182).
 - Access-Aggregation-Circuit-ID-Binary—Identifier that represents the outer VLAN tag that the access node inserts on upstream traffic; corresponds to the Juniper Networks Acc-Aggr-Cir-Id-Bin VSA (26-111).
- The name of the physical interface facing the subscriber. This name derives from the local mapping of an ANCP neighbor to the corresponding subscriber-facing access port.

The Access-Aggregation-Circuit-ID-Binary attribute and the access-facing interface name together provide information equivalent to that used for conventional autosensed VLAN detection.

ANCP Port Down messages indicate that the subscriber access loop is not present or at least is no longer operational. This message triggers the automatic destruction of the dynamic VLAN, regardless of the value of any other ANCP line attribute.

VLAN logical interfaces are created in the default routing-instance unless a nondefault routing instance is provided by local authorization (domain map) or external authorization (RADIUS). Multiple routing instances are required when both access-provider-owned and wholesaled connections are supported at the same time. One routing instance is required for the access provider's own subscribers. An additional routing instance is required for each retail NSP. Consequently, the routing-instance has to be specified when the VLAN is authorized. The RADIUS-based VLAN authorization process determines whether the subscriber access-loop identified by the attributes in the Port Up message is wholesaled to a partner NSP—and therefore maintained as a unique routing-instance—or managed as a subscriber owned by the access provider.

RADIUS Authorization for ANCP-Triggered VLANs

When a subscriber logs in, the Access-Request message that is sent to the RADIUS server includes a username and optionally a password generated locally on the router. You can configure the router to create a unique username with the value of ANCP TLVs— Access-Loop-Circuit-ID, Access-Loop-Remote-Id, or both—as received in the ANCP Port Up message from the access node. Alternatively, if you configure the router to convey ANCP-sourced access loop attributes as Juniper Networks VSAs—in this case Acc-Loop-Cir-Id (26-110) and Acc-Loop-Remote-Id (26-182)—then the Access-Request message includes sufficient unique access line information for the RADIUS server to determine whether the access loop is wholesaled to a retailer or retained for the wholesaler.

The RADIUS server responds to the Access-Request with one of the following messages:

- Access-Accept—In this case, the VLAN triggered by the Port Up message is wholesaled to a retail NSP. Authorization is similar to that for PPPoE sessions. The Access-Accept includes the Virtual-Router VSA (26-1) with a value that corresponds to the NSP's unique, nondefault routing instance. The message can optionally include client services, such as attributes for parameterized CoS, firewall filters and policies for the logical interface, and Layer 2 service activations.
- Access-Reject—In this case, either the VLAN triggered by the Port Up message is one of the wholesaler's own subscribers or RADIUS refuses to grant access to the network. In either case, the VLAN entry is removed from the ANCP SDB. Unless a Port Down message is received first, the router ignores subsequent Port Up messages for this subscriber. However, conventional dynamic stacked VLAN autosensing may be initiated by access protocol negotiation, such as PPPoE.

Instantiation of an ANCP-Triggered, Autosensed, Dynamic VLAN

When the RADIUS server returns an Access-Accept message, the dynamic profile assigned to the autosensed VLAN range is instantiated with the following results:

1. The dynamic VLAN logical interface that represents the Layer 2 wholesale service within the NSP's unique routing instance is created.
2. A core-facing physical interface is selected by a weighted load distribution method from the set of eligible interfaces assigned to the NSP's routing instance. A physical interface is eligible when it is operationally up and has at least one VLAN tag that is available for assignment.
3. The access-facing, autosensed outer VLAN tag is mapped to a unique inner VLAN tag. The outer VLAN tag is derived from the Access-Aggregation-Circuit-ID-Binary TLV carried in the ANCP Port Up message. The inner VLAN tag is allocated from the VLAN range configured for the core-facing physical interface.
4. The inner VLAN tag is swapped with (replaces) the outer VLAN tag when the subscriber traffic is tunneled to the NSP. In the dynamic profile, the inner VLAN tag is provided by the predefined variable, **\$junos-inner-vlan-map-id**.

5. The autosensed outer VLAN tag is swapped with the inner VLAN tag when downstream packets from the NSP (which include the allocated inner VLAN tag) are forwarded to the subscriber.

You can configure each core-facing physical interface with a range of up to 4094 VLAN IDs. The inner VLAN swap range is assigned to the physical interface locally. This means that inner VLAN ranges for different physical interfaces can overlap each other completely, partially, or not at all.

6. Optionally, before the subscriber packets are forwarded to an NSP, the outer VLAN tag protocol identifier (TPID) in the subscriber packets can be swapped with a TPID to meet the requirements of an individual NSP. In this case, the original value is swapped with the NSP TPID for packets forwarded to the subscriber.
7. An additional VLAN tag, the Trunk VLAN ID, is used internally to identify the provisioned core-facing physical interface so that the subscriber traffic can be tunneled to the allocated interface. In the dynamic profile, this ID is provided by the predefined variable, **\$junos-vlan-map-id**. This identifier differentiates among multiple core-facing trunk physical interfaces for the same NSP.
8. Any client services, such as CoS or firewall filters, are applied to the subscriber session. These services are optionally specified in the RADIUS configuration and conveyed in the RADIUS message as Juniper Networks VSAs.
9. The VLAN session is activated after the logical interface is created and configured for the dynamic VLAN session. Session activation initiates a RADIUS Accounting-Start message. Any services that were received from RADIUS during authorization are now activated.
10. After the dynamic VLAN has been created, subsequent ANCP Port Up messages do not cause a re-authorization of the dynamic VLAN session. Instead, when the ANCP agent receives another Port Up message for the access loop, it updates the ANCP SDB with any changes in ANCP attributes.

Weighted Load Balancing for Subscriber Sessions over Eligible Core-Facing Physical Interfaces

The router uses weighted load distribution instead of round-robin distribution to assign Layer 2 wholesale subscriber sessions across multiple core-facing physical interfaces according to the weight of the interface. The weight of an interface correlates with the number of VLAN tags available from the aggregate inner (core-facing) VLAN ID swap ranges on the interface.

How you configure the inner VLAN ID swap ranges determines the relative weights of the interfaces:

- The interface with the highest number of available inner VLAN ID tags has the highest weight.
- The interface with the next-highest number of tags has the next-highest weight, and so on.
- The interface with the lowest number of available tags has the lowest weight.

Using the available inner VLAN ID tags from the swap ranges rather than the aggregate total VLAN tags means that the relative weights of the interfaces are more dynamic. The weighted load distribution mechanism can respond more quickly to subscriber logouts, migration of subscriber ownership from wholesaler to retailer and retailer to wholesaler, core-facing physical interface state transitions (including movement to the remaining eligible core-facing interfaces when an interface state transitions from Up to Down), and failures of any of the core-facing physical interfaces. When an interface recovers (transitions from Down to Up), weighted load distribution generally favors this interface for either pending sessions or for new Layer 2 wholesale sessions that subsequently occur.

NOTE: Core-facing physical interface selection and session distribution are probability based; the load is not strictly distributed according to weight.

With weighted load distribution the router selects interfaces randomly, but the sessions are distributed across interfaces proportionally in relationship to the weight of the interfaces. The router generates a random number within a range equal to the aggregate total of all available inner VLAN ID tags from the swap ranges of all the core-facing physical interfaces. The router then associates part of the range—a pool of numbers—with each interface proportional to the interface's weight. An interface with a higher weight is associated with a greater portion of the range—a larger pool—than an interface with a lower weight. An interface is selected when the random number is in its associated pool of numbers. The random number is more likely, on average, to be in a larger pool, so an interface with a higher weight (larger pool) is more likely to be selected than an interface with a lower weight (smaller pool).

For example, consider two core-facing physical interfaces, IFD1 and IFD2. Based on the inner VLAN ID swap ranges configured for these two interfaces, IFD1 has 1000 available VLAN tags and IFD2 has only 500 available tags. The subscriber sessions are randomly distributed across the two interfaces based on their relative weights; IFD1 has a higher weight than IFD2. Because IFD1 has twice as many available tags as IFD2, the pool of numbers associated with IFD1 is also twice as many as for IFD2. The random number generated by the router is twice as likely to be in the pool for IFD1 than for IFD2. Consequently, IFD1 is favored 2:1 over IFD2, and subscriber sessions are twice as likely to be assigned to IFD1 as IFD2.

RADIUS Interim Accounting Updates

Interim accounting reports sent to AAA for out-of-band triggered, autosensed dynamic VLANs are supported in the same manner as for conventional autosensed, dynamic, authorized VLANs or client sessions (such as PPPoE sessions). The ANCP agent sends a notification to AAA when it receives an update from the access node. By default, AAA only reports the update to the RADIUS server at the configured interval.

You can configure the ANCP agent so that when it notifies AAA, an interim update Accounting-Request message is immediately sent to the RADIUS server. Immediate interim accounting updates can be sent for an ANCP-triggered dynamic VLAN session only when a change occurs in certain key ANCP attributes for the associated access line that can influence system behavior. To prevent an additional load on the RADIUS server for changes to less critical ANCP attributes, changes to any other ANCP attributes do not trigger

immediate accounting-interim-update messages. Instead, those changes are reported in the next scheduled Accounting-Interim-Update message.

Immediate interim accounting updates can be sent for changes to any of the following ANCP attributes for an existing session that corresponds to the access line (based on the Access-Loop-Circuit-ID TLV):

- **Actual-Net-Data-Rate-Upstream**—When the calculated (adjusted) upstream rate results in a change to this attribute, the accounting message reports the attribute in the Juniper Networks Act-Data-Rate-Up VSA (26-113). The calculated speed change is reported in the Upstream-Calculated-QoS-Rate VSA (26-142).
- **Actual-Net-Data-Rate-Downstream**—When the calculated (adjusted) downstream rate results in a change to this attribute, the accounting message reports the attribute in the Juniper Networks Act-Data-Rate-Dn VSA (26-114). The calculated speed change is reported in the Downstream-Calculated-QoS-Rate VSA (26-141).

When the **ancp-speed-change-immediate-update** statement is configured at the **[edit access profile profile-name accounting]** hierarchy level, RADIUS immediate interim accounting updates are sent for changes to the Actual-Net-Data-Rate-Upstream and Actual-Net-Data-Rate-Downstream TLVs.

When in addition the **auto-configure-trigger interface interface-name** statement is configured at the **[edit protocols ancp neighbor ip-address]** hierarchy level, immediate interim accounting updates are also sent for changes to the Access-Loop-Remote-ID and Access-Aggregation-Circuit-ID-Binary TLVs.

For more information about RADIUS immediate interim accounting updates, see *Configuring Immediate Interim Accounting Updates to RADIUS in Response to ANCP Notifications*.

Removal of the Layer 2 Wholesale Service

Any of the following events can remove the logical interface for the dynamic VLAN that represents the access service provided by the Layer 2 wholesaler:

- The receipt of an ANCP Port Down message for the corresponding access loop. The same ANCP attributes that initiate dynamic VLAN creation also initiate dynamic VLAN destruction.

No action is taken for an ANCP Port Down message for which any of the following is true:

- No corresponding subscriber session exists.
- A corresponding subscriber session is present, but is in the process of being deleted.
- The message refers to a conventional autosensed session (which is removed by normal protocol processing).
- An explicit reset of the connection between the ANCP agent and the access node, which triggers a mass logout of all affected dynamic VLAN sessions that support the wholesaled Layer 2 access connections. Sessions for the wholesaler's own subscribers are not affected.

- The deletion or transition to an operational down state of the subscriber-facing physical interface or the core-facing physical interface.
- The loss of adjacency between the neighbor and the ANCP agent.
- The issuance of the **clear auto-configuration interfaces** command to log out the VLAN or the **clear ancp access-loop** command to force a subscriber reset.
- The receipt of a RADIUS-initiated disconnect message.

Any of these events also deactivates the subscriber session to prevent future service activations and issues a RADIUS Accounting-Stop message for related services and for the dynamic VLAN subscriber session. The dynamic profile is then de-instantiated to remove first the dynamic VLAN logical interface and then the corresponding session entry in the VLAN SDB.

You can monitor the number of Layer 2 cross-connected subscriber sessions per port. Use the **show subscribers summary port extensive** command to display the number of subscribers per port by client type (VLAN-OOB) and connection type (Corss-connected). Additionally, the object ID `jnxSubscriberPortL2CrossConnectCounter` in the `jnxSubscriberPortCountTable` in the Juniper Networks enterprise-specific Subscriber MIB displays the number of Layer 2 cross-connected subscriber sessions on ports that have active sessions.

Interactions Between In-Band and Out-of-Band VLAN Autosensing

The ANCP-triggered Layer 2 wholesale implementation accommodates both subscribers wholesaled to a retailer and subscribers belonging to the wholesaler. Any subscriber session detected on the access-facing physical interface can be one or the other. This means that an overlap exists between the outer tag range for the out-of-band autosensed VLANs and that for in-band, autosensed, stacked VLANs.

Both a PPPoE session and a wholesaled session might be attempted for the same access loop. To avoid the undesirable load on the router and the RADIUS server that can ensue when that happens, the order of session negotiation is determined by the order in which packets (PPPoE PADI or ANCP Port Up message) are received for the same access-facing physical interface and VLAN outer tag.

NOTE: The following sequences assume that the **remove-when-no-subscribers** statement is included at the **[edit interfaces interface-name auto-configure]** hierarchy level for the access-facing physical interface.

The following sequence of events occurs when a PPPoE PADI packet is received on an in-band control channel before an ANCP Port Up message is received on an out-of-band control channel, for the same access loop:

1. The PADI receipt triggers creation of a dynamic stacked VLAN logical interface. PPPoE and PPP negotiation are in progress.
2. The ANCP Port-Up message is received for the access loop. Because the corresponding in-band VLAN logical interface already exists for the same access-facing physical interface and outer VLAN tag, the ANCP agent simply stores the ANCP access line attributes and the name of the physical interface in the session database. The agent takes no other action for the message as long as the PPP session (PPP logical interface and the underlying dynamic VLAN logical interface) is maintained.
3. PPP negotiation terminates due to authentication failure (RADIUS Access-Reject response) or a normal logout, which triggers clean-up of the PPP session and removal of the PPP logical interface.
4. Because the **remove-when-no-subscribers** statement is configured, deletion of the PPP logical interface results in deletion of the dynamic stacked VLAN.
5. The next event depends on whether authorization of the ANCP Port Up message has been attempted before.
 - If authorization was not previously attempted:
 - a. A VLAN-OOB SDB session is created and authorization of the access-loop is initiated.
 - b. All exceptioned PPPoE PADI packets received by in-band VLAN auto-sensing are ignored until RADIUS responds to the authorization request.
 - c. The authorization result determines what happens next:
 - If the authorization succeeds (RADIUS returns an Access-Accept message), then a dynamic Layer 2 wholesale logical interface is created within the retailer NSP's routing-instance.
 - If the authorization fails (RADIUS returns an Access-Reject message), then the VLAN-OOB session is cleaned up. Processing resumes for any exceptioned PPPoE PADI packets that are subsequently received by in-band VLAN autosensing.
 - If authorization was previously attempted, then no action is taken because the failure of the PPP session negotiation is assumed to be a login failure outside the Layer2 wholesale context. This behavior prevents unnecessary authorization in response to the ANCP Port-Up message every time a PPPoE session terminates and cleans up from a normal subscriber logout.

The following sequence of events occurs when an ANCP Port Up message is received on an out-of-band control channel before a PPPoE PADI packet for an access loop is received on an in-band control channel, both for the same access loop:

1. Receipt of the ANCP Port Up message triggers authorization of the access loop.
2. A PPPoE PADI packet is received. The packet is ignored because authorization is already in progress for the same access-facing physical interface and outer VLAN tag.
3. The authorization result determines what happens next:
 - If authorization succeeds (RADIUS returns an Access-Accept message)—represented by a VLAN-OOB session in the SDB—then dynamic creation of the VLAN logical interface is initiated for a Layer 2 wholesale session. When the interface is created, subsequent PPPoE PADI packets detected by in-band VLAN autosensing are ignored and no longer exceptioned.
 - If authorization fails (RADIUS returns an Access-Reject message), the VLAN-OOB session is cleaned up.
 - a. Receipt of a subsequent PPPoE PADI packet initiates creation of a dynamic stacked VLAN.
 - b. PPPoE and PPP negotiation takes place and events proceed as usual for a conventional, dynamic autosensed VLAN.

Migration of Subscriber Ownership from Wholesaler to Retailer

The wholesaler-owned subscribers are connected by means of dynamic PPPoE interfaces over dynamic VLANs. For each subscriber, the PPPoE session must be disconnected and the corresponding PPP logical interface deleted before ANCP Port Up messages for the same underlying physical interface and VLAN outer tag can serve as out-of-band triggers for dynamic VLAN autosensing.

One approach to migrating from wholesale to retail ownership is to do the following:

1. Update the RADIUS server database so that subscriber authentication for the relevant access lines results in a RADIUS Access-Reject response. This causes subsequent attempts to negotiate PPPoE for the access line to fail.
2. Initiate logout of the dynamic PPPoE sessions; for example, by issuing a RADIUS-initiated disconnect. This triggers cleanup of the PPPoE logical interface and associated services, which includes issuing RADIUS Accounting-Stop messages for the session and active services, as well as removing the dynamic PPPoE logical interface.

If the migration requires swapping out the current CPE device for another, and the PPPoE session is not otherwise gracefully logged out, then turning off the CPE results in a PPP keepalive failure on the router and triggers session logout.

3. Remove the underlying dynamic VLAN logical interface. This occurs automatically when the **remove-when-no-subscribers** statement is included at the **[edit interfaces *interface-name* auto-configure]** hierarchy level for the access-facing physical interface. Otherwise, issue the **clear auto-configuration interfaces *interface-name*** command to remove the dynamic VLAN logical interface.
4. Trigger a Port Up notification to initiate dynamic VLAN detection, authorization, and creation by one of the following methods:
 - Power cycle the CPE, with a sufficient delay between turning off and turning back on the device so that a Port Down message is sent followed by a Port Up message and the router is given enough time to detect a keepalive failure indicating loss of the session.
 - Issue a **clear ancp access-loop** command.
 - Issue a **request ancp oam port-up** command.

Migration of Subscriber Ownership from Retailer to Wholesaler

One approach to migrating from retail to wholesale ownership is to do the following:

1. Update the RADIUS server database so that dynamic VLAN authorization for the relevant access lines results in a RADIUS Access-Reject response. Doing this causes subsequent ANCP Port Up messages to be ignored.
2. Initiate logout of the dynamic VLAN sessions supporting the wholesale access service; for example, by issuing a RADIUS-initiated disconnect. Doing this triggers cleanup of the session, which includes issuing RADIUS Accounting-Stop messages for the session, removal of the dynamic VLAN logical interface and active services, and freeing the allocated inner VLAN tag associated with the core-facing physical interface for assignment to a different Layer 2 wholesale subscriber session.

If the migration requires swapping out the current CPE device for another, then turning off the CPE results in an ANCP Port Down message that triggers session logout and cleanup.

3. Allow subscribers to connect to the wholesaler's network using conventional dynamic VLAN autosensing followed by PPPoE and PPP negotiation and creation of PPP logical interfaces.

Migration of Subscriber Ownership Between Retailers

Typically, you migrate access between NSP retailers by triggering a restart of the existing dynamic VLAN session. The restart initiates a logout from the session followed by immediate dynamic VLAN detection, authorization, and creation within the routing-instance corresponding to the new NSP. A restart is a logical Port Down and Port Up sequence for the VLAN's corresponding access loop. You can use any of the following methods to restart a given dynamic VLAN logical interface:

- Initiate a RADIUS Disconnect-Request message or configure your RADIUS server to send the message. The message must have the Acct-Terminate-Cause RADIUS attribute (49) set to a value of 16 (callback). This cause is processed as a disconnect (logout) followed immediately by a reconnect (login) only for dynamic VLANs created by an ANCP Port Up message. Other clients respond to this value with only a disconnect.
- Include the **reconnect** option when you log out subscribers with the **clear network-access aaa subscriber** command. You can specify subscribers by either the session identifier or the username. This option attempts to reconnect a cleared session as a Layer 2 wholesale session when the subscriber session has been fully logged out. This behavior is equivalent to issuing a RADIUS-initiated disconnect that is configured for reconnect; that is, when the message includes Acct-Terminate-Cause (RADIUS attribute 49) with a value of callback (16).
- Trigger a Port Down message followed by a Port UP message by one of the following methods:
 - Power cycle the CPE, with a sufficient delay between turning off and turning back on the device so that a Port Down message is sent followed by a Port Up message and the router is given enough time to detect a keepalive failure indicating loss of the session.
 - Issue a **clear ancp access-loop** command.

Modification of the Access Line Identifier or Port VLAN Identifier

When the line identifier or port VLAN identifier for an access loop is modified, the access node reports the change in a Port Up message to the ANCP agent. The message conveys the line ID in the Access-Loop-Remote-ID TLV and the port VLAN ID in the Access-Aggregation-Circuit-ID-Binary TLV.

The access node should send a Port Down message for the access loop before it modifies any of the identification attributes for an existing session. The Port Down message triggers clean up of the corresponding Layer 2 wholesale session. If the access node does not send a Port Down in this case, then the following behavior has the same effect as inserting the Port Down message that the access node failed to send:

- For a line ID change, the ANCP agent treats the reconfiguration as a logical Port Down message for the access line identified by the previous Access-Loop-Remote-Id, followed by a Port Up message for the access line identified by the new Access-Loop-Remote-Id.

- For a port VLAN ID change, the ANCP agent treats the reconfiguration as a logical Port Down message for the access line identified by the previous Access-Aggregation-Circuit-Id-Binary, followed by a Port Up message for the access line identified by the new Access-Aggregation-Circuit-Id-Binary.

In either case, the ANCP agent notifies the autoconfiguration daemon (autoconfd), which takes the following actions:

1. Logs out the corresponding Layer 2 wholesale session.
2. Sends a RADIUS Accounting-Stop message with the final statistics for the associated service sessions and client session.
3. Removes the dynamic VLAN logical interface.
4. Attempts to reestablish the Layer 2 wholesale session by means of a login sequence, including authentication, creation of the dynamic VLAN logical interface, activation of any services, and if successful, sending RADIUS Accounting-Start messages for the service and client sessions.

You must manually log out any PPPoE session corresponding to the access line with the previous identifiers, even if the access node sends the appropriate Port Down message when the values change.

NOTE: In the case of a change in the port VLAN ID only, autoconfd takes no action to reinitiate the session when a dynamic stacked VLAN or a Layer 2 wholesale session exists with the same access-facing physical interface and outer VLAN tag. You must manually intervene in this case, such as by issuing a **request ancp oam port-up** command to initiate the creation of the Layer 2 wholesale session.

BEST PRACTICE: Because an existing session is not automatically logged out, we recommend that the network operator disconnect the session—for example, by issuing a RADIUS Disconnect-Request message—before modifying any of the access line attributes. Depending on subsequent subscriber login and successful negotiation, a new session with the new identifier may then be created as usual.

Disconnecting PPPoE Sessions and Automatically Attempting Reconnection as Layer 2 Wholesale Sessions

You can use RADIUS-initiated disconnect messages to disconnect and log out existing PPPoE sessions and attempt to reestablish them as Layer 2 wholesale sessions. Use Access-Reject messages to prevent PPPoE subscriber access and attempt a reconnect as a Layer 2 wholesale session. Use this feature when you want to migrate sessions from PPPoE to Layer 2 wholesale. Both the RADIUS-initiated disconnect and Access-Reject message must include Acct-Terminate-Cause (RADIUS attribute 49) with a value of

callback (16); callback causes the reconnect attempt. The **remove-when-no-subscribers** statement must be configured on the underlying physical interface.

1. The initial behavior for the messages is the following:

- Access-Reject message—When a PPPoE PADI is received and a new PPPoE session is requested, RADIUS responds to the Access-Request message with an Access-Reject message. The session is rejected, fully logged out, and the underlying dynamic VLAN logical interface is removed.
- RADIUS-initiated disconnect message—When a RADIUS-initiated disconnect message is received for an existing PPPoE session, the dynamic PPPoE session is logged out and the underlying dynamic VLAN logical interface is removed.

2. The next action is the same for both messages:

- If an ANCP Port Up message has been received for the corresponding access line, the router attempts to authorize the access line and create a dynamic Layer 2 wholesale VLAN logical interface in place of the underlying dynamic VLAN logical interface that was removed.
- If a Port Up message has not been received, then this action is deferred until the message is received.
- If a PPPoE PADI is received before an ANCP Port Up message, RADIUS responds to the Access-Request for a new PPPoE session with an Access-Reject message. The session is rejected, fully logged out, and the underlying dynamic VLAN logical interface is removed.

If the RADIUS-initiated disconnect or Access-Reject message is received for a non-PPPoE session, such as DHCP, the session is disconnected, but the reconnect request is ignored. No attempt is made to establish a Layer 2 wholesale session.

If the RADIUS-initiated disconnect does not include Acct-Terminate-Cause with a value of callback, PPPoE renegotiation after the disconnect can succeed, but if an ANCP Port Up message is received for the access line before a PPPoE session is established, then a Layer 2 wholesale session is attempted.

As an alternative to the RADIUS-initiated disconnect, you can manually log out the PPPoE session with the **clear network-access aaa subscriber** command. Specify the subscriber by either username or session ID. When you include the **reconnect** option, it attempts to reconnect the cleared session as a Layer 2 wholesale session when the subscriber session has been fully logged out.

Consequences of a State Transition in the Access-Facing Physical Interface

The following behavior results when the access-facing physical interface state transitions from Up to Down:

- Conventional in-band VLAN autosensing stops for the interface.
- ANCP-sourced Port Up messages for the interface are ignored. Action on new or unprocessed Port Up messages is deferred until the interface transitions to the Up state. If the ANCP connection is in band with the subscriber traffic on the interface, then all ANCP traffic stops; if the outage lasts long enough, the ANCP adjacency is lost.

- All Layer 2 wholesale sessions that are assigned to the interface are treated as if the ANCP agent received a Port Down message for the corresponding access line. Each session is subject to being logged out. Whether a session is logged out depends on the ANCP adjacency loss hold timer. The timer starts running when the ANCP agent detects the state transition to Down. The subscriber continues using the original session if all three of the following occur before the timer expires:

1. The physical interface comes back up.
2. The ANCP adjacency is restored.
3. A Port Up message is received on the interface.

Otherwise, autoconfd takes the following actions:

1. Logs out the session.
2. Sends a RADIUS Accounting-Stop message with the final statistics for the associated service sessions and client session.
3. Removes the dynamic VLAN logical interface.

These sessions can be reestablished when the physical interface recovers, unless an ANCP Port Down message is received.

- The autoconfiguration daemon does not automatically delete dynamic, autosensed VLAN logical interfaces. The interfaces for the ANCP-triggered Layer 2 wholesale VLANs are maintained because the assumption is that an outage is short-lived. If the outage is not short-lived, then a subsequent Port Down message brings down the session and removes the interface.

For conventional autosensed dynamic VLANs, the interface is removed only when the **remove-when-no-subscribers** statement is configured on the access-facing physical interface and all references to the VLAN logical interface from a higher logical interface or session are removed. This mechanism does not apply to the ANCP-triggered Layer 2 wholesale VLANs because they do not have upper session references.

The following behavior results when the access-facing physical interface state transitions from Down to Up:

1. Conventional in-band VLAN autosensing resumes for the interface. PPPoE sessions owned by the access provider that were logged out due to the transition from Up to Down can renegotiate and undergo a full login sequence.
2. Appropriate actions are taken for all ANCP Port Up messages for the interface, including messages that were deferred because of the previous Down state for the interface. If the ANCP connection is in band with the subscriber traffic, then all ANCP traffic resumes.
3. Forwarding resumes for any dynamic, autosensed VLAN logical interfaces that were not deleted when the interface went down.

Deletion of an access-facing physical interface triggers logout and removal of all upper dynamic VLAN logical interfaces and their corresponding sessions.

Consequences of a State Transition from Up to Down in the Core-Facing Physical Interface

The following behavior results when the core-facing physical interface state transitions from Up to Down:

- The core-facing physical interface is no longer eligible for assigning new or pending access lines in this routing instance as based on the original RADIUS authorization.
- All Layer 2 wholesale sessions that are assigned to the interface are treated as if the ANCP agent received a Port Down/Port Up message sequence for the corresponding access line. Each session is subject to being logged out. Whether a session is logged out depends on the ANCP adjacency loss hold timer. The timer starts running when the ANCP agent detects the state transition to Down. The subscriber continues using the original session if all three of the following occur before the timer expires:
 1. The physical interface comes back up.
 2. The ANCP adjacency is restored.
 3. A Port Up message is received on the interface.

Otherwise, autoconfd takes the following actions:

1. Logs out the session.
 2. Removes the session entry from the SDB.
 3. Sends a RADIUS Accounting-Stop message with the final statistics for the associated service sessions and client session.
 4. Removes the dynamic VLAN logical interface.
- Next, autoconfd attempts to migrate the sessions to available connections on any remaining eligible core-facing physical interfaces that are assigned to the same routing instance:
 1. The original request is placed on a retry queue.
 2. A login sequence is attempted for each session, including authentication, creation of dynamic VLAN logical interfaces, activation of any services, and sending RADIUS Accounting-Start messages for the service and client sessions.
 - If the login sequence is successful, then the request is removed from the retry queue.
 - If the login fails, then the session is logged out, the session entry is removed from the SDB, and the corresponding access line is set to a pending state.

When the available connections are all used—when there are no more available VLAN tags from the configured inner VLAN ID swap ranges—as a result of successful reconnections, no attempt is made to connect any remaining Layer 2 wholesale sessions. Although authentication can succeed, the creation of dynamic VLAN logical interfaces fails during profile instantiation. In this case, the session is out, the session entry is removed from the SDB, and the corresponding access line is set to a pending state.

- The pending access lines that represent these non-migrated sessions can be reestablished if the interface recovers or if additional VLAN connections become available; for example, by a configuration change

that either increases the inner VLAN ID swap range for one or more remaining core-facing physical interfaces or adds new core-facing physical interfaces. However, if the ANCP agent receives a Port Down message for a pending access line, the corresponding session is not reestablished.

You can use the **show auto-configuration out-of-band pending** command to display a count of pending access lines per routing instance.

NOTE:

In addition to core-facing physical interface state transitions from Up to Down, these behaviors also apply in the following circumstances:

- A core-facing physical interface is deleted.
- More Layer 2 wholesale sessions are assigned to a routing instance than can be accommodated by the inner VLAN ID swap range configured on the interface assigned to the routing instance.

BEST PRACTICE: We recommend that you use aggregated Ethernet for the core-facing physical interfaces to provide link protection, bandwidth aggregation, or both.

Consequences of a State Transition from Down to Up in the Core-Facing Physical Interface

The following behavior results when the core-facing physical interface state transitions from Down to Up:

- The physical interface is now eligible to assign new Layer 2 wholesale subscriber sessions.
- The ANCP agent notifies the autoconfiguration daemon (autoconfd), which attempts to reestablish the Layer 2 wholesale sessions that correspond to pending access line by initiating a conventional login sequence. This sequence includes authentication, creation of dynamic VLAN logical interfaces, activation of any services, and sending RADIUS Accounting-Start messages for the service and client sessions.
- Pending sessions continue to be reestablished until none are left or an error occurs, typically due to exhaustion of inner VLAN tags from the swap ranges on the interface. In the latter case, the sessions are logged out, the session entry is removed from the SDB, and the access line is set to a pending state.

You can use the **show auto-configuration out-of-band pending** command to display a count of pending access lines per routing instance.

These behaviors also occur in the following cases:

- Additional VLAN connection resources become available, by a configuration change that either increases the inner VLAN ID swap range for one or more remaining core-facing physical interfaces or adds new core-facing physical interfaces. The newly added physical interface must be in the Up state to assume any Layer 2 wholesale sessions.

- A RADIUS-initiated disconnect is received for an existing Layer 2 wholesale session assigned to this routing instance is logged out (disconnect only). For a disconnect with a reconnect qualifier, the affected session is given preference to reconnect over pending access lines.
- You issue the **request auto-configuration reconnect-pending**, **clear ancp access-loop**, or **request ancp oam port-up** command.

Loss of ANCP TCP Adjacency

The ANCP agent can lose its TCP adjacency with a neighbor in any of the following circumstances:

- The access node renegotiates the connection; for example, as a result of losing synchronization. The renegotiation triggers the local state to change from established to not established. The state transitions back to established when the session is successfully renegotiated.
- An end-of-file (EOF) message is received on the socket indicating the adjacency is closed. This can result when the ANCP configuration is deleted on the access node.
- An adjacency keepalive failure occurs. When no response is received for three consecutive polls, the adjacency is declared to be lost.

The ANCP agent treats the loss of adjacency as if it has received a Port Down message for each access loop represented by the ANCP connection. The agent notifies autoconfd, which takes the following actions:

- Logs out all Layer 2 wholesale sessions that were triggered by this ANCP connection.
- Sends a RADIUS Accounting-Stop message with the final statistics for the associated service sessions and client session.
- Removes the dynamic VLAN logical interface.

If the assigned access-facing or core-facing physical interface is in the Down state, any pending sessions that were triggered by this ANCP connection cannot be reestablished when the interface recovers to the Up state.

Dynamic, conventionally auto-sensed VLAN logical interfaces, such as those supporting PPPoE sessions, are not affected by the TCP adjacency loss.

If the adjacency is reestablished, the expected behavior is a complete replay of Port Down and Port Up messages for all associated configured access lines. The Layer 2 wholesale sessions for which the ANCP agent receives Port Up messages are reestablished.

You can mitigate the effects of short-term adjacency losses by configuring an adjacency loss hold time. The timer starts when adjacency is lost. Even though the adjacency is lost, established sessions are maintained while the timer runs unless a Port Down message is received for the corresponding access line.

Any access line that for which the ANCP agent has not received a Port Up message by the time the timer expires is treated as though the agent has received a Port Down message for the line. The ANCP Agent notifies autoconfd, which takes the following actions:

- Logs out all Layer 2 wholesale sessions that correspond to the access line.
- Sends a RADIUS Accounting-Stop message with the final statistics for the associated service sessions and client session.
- Removes the dynamic VLAN logical interface.

Port Up messages received after the timer expires repopulate the SDB access line table and reestablish the Layer 2 wholesale sessions

The adjacency loss hold timer serves the following purposes:

- Dampens the effect of adjacency loss of short duration thereby maintaining existing Layer 2 wholesale sessions.
- Detects the removal of an access line configuration on the access node. For example, in some circumstances you may want to remove the configuration of an access line on a neighbor. You first disconnect the ANCP session between a neighbor and the BNG, remove the configuration on the neighbor, and then restore the ANCP connection with the BNG. The neighbor does not issue a Port Down message. If the adjacency-loss hold-timer is configured, the ANCP agent detects an access line for which it has not received a Port Up or Port Down message, and consequently triggers logout and removal of the corresponding Layer 2 wholesale session.

NOTE: When you deactivate the ANCP protocol, the router does not perform a commit check to determine whether any ANCP or L2-BSA subscribers are present (active or inactive). Any subscribers that are active at the time of deactivation remain active.

Release History Table

Release	Description
16.1R4	Starting in Junos OS Release 16.1R4, you can configure the ANCP agent to trigger the creation of an autosensed VLAN when the ANCP agent receives a Port Up message where the DSL-Line-State attribute has a value of Showtime.

RELATED DOCUMENTATION

[Configuring Out-of-Band ANCP Messages to Trigger Dynamic VLAN Instantiation](#) | 141

[Configuring the ANCP Agent for ANCP-Triggered, Autosensed Dynamic VLANs](#) | 138

<i>Configuring the ANCP Agent</i>
<i>ANCP and the ANCP Agent Overview</i>
<i>Junos OS Predefined Variables</i>
<i>Juniper Networks VSAs Supported by the AAA Service Framework</i>
<i>AAA Access Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS</i>
<i>AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS</i>

Configuring ANCP-Triggered Layer 2 Wholesale Services

IN THIS CHAPTER

- Configuring ANCP Neighbors | 136
- Configuring the ANCP Agent for ANCP-Triggered, Autosensed Dynamic VLANs | 138
- Configuring a Username for Authentication of Out-of-Band Triggered Dynamic VLANs | 139
- Configuring Out-of-Band ANCP Messages to Trigger Dynamic VLAN Instantiation | 141
- Triggering ANCP OAM to Simulate ANCP Port Down and Port Up Messages | 142
- Configuring the ANCP Agent to Dampen the Effects of Short-Term Adjacency Losses | 144
- Reestablishing Pending Access Line Sessions for Layer 2 Wholesale | 145
- Configuring Multiple Non-Overlapping VLAN Ranges for Core-Facing Physical Interfaces | 146
- Clearing ANCP Access Loops | 147

Configuring ANCP Neighbors

You must configure each neighboring access node that you want the ANCP agent to monitor and potentially shape traffic for. Some neighbor settings override globally configured values.

To configure an ANCP neighbor:

1. Specify the IP address of the neighbor.

```
[edit protocols ancp]
user@host# set neighbor 203.0.113.234
```

2. (Optional) Configure the neighbor to operate in a backward-compatible mode when it does not support the current IETF standard and the backward-compatible mode is not configured globally.

```
[edit protocols ancp neighbor 203.0.113.234]
user@host# set pre-ietf-mode
```

3. (Optional) Override the globally configured backward-compatible mode when the neighbor supports the current IETF standard.

```
[edit protocols ancp neighbor 203.0.113.234]  
user@host# set ietf-mode
```

4. (Optional) Configure the interval in seconds between ANCP adjacency messages exchanged with this neighbor.

```
[edit protocols ancp neighbor 203.0.113.234]  
user@host# set adjacency-timer 20
```

5. (Optional) Specify the maximum number of discovery table entries that are accepted from this neighbor.

```
[edit protocols ancp neighbor 203.0.113.234]  
user@host# set maximum-discovery-table-entries 10000
```

6. (Optional) Enable out-of-band ANCP triggering of autosensed, dynamic VLANs on the physical interface.

```
[edit protocols ancp neighbor 203.0.113.234]  
user@host# set auto-configure-trigger interface ge-1/0/0
```

7. (Optional) Configure how long the ANCP agent maintains a Layer 2 wholesale session when an adjacency loss occurs.

```
[edit protocols ancp neighbor 203.0.113.234]  
user@host# set adjacency-loss-hold-time 10
```

Configuring the ANCP Agent for ANCP-Triggered, Autosensed Dynamic VLANs

Starting in Junos OS Release 16.1R4, you can configure the ANCP agent to associate a neighbor with an access-facing physical interface for the creation of autosensed dynamic VLANs on the interface. When the ANCP agent receives a Port Up message from the neighbor, it triggers notification to the autoconfd daemon to initiate the detection, authorization, and creation of dynamic VLANs. Receipt of an out-of-band ANCP Port Down message triggers notification to the autoconfd daemon to initiate the destruction of an existing VLAN on the interface.

NOTE: The following physical interface types are supported: aggregated Ethernet (ae), Gigabit Ethernet (ge), 10-Gigabit Ethernet (xe), 100-Gigabit Ethernet (et), demux, and pseudowire (ps). The ps interface type was added in Junos OS Release 19.3R1.

This configuration assumes the following:

- The dynamic profile is configured to instantiate a dynamic VLAN when notified by the ANCP agent that it has received an out-of-band ANCP Port Up message.
- The RADIUS authentication server is properly configured to authorize the VLANs and apply services as needed.
- The ANCP agent is configured to initiate interim accounting updates (which also enables immediate interim accounting updates) in response to information received in Port Up messages.

To map a neighbor to a physical interface for autosensed dynamic VLANs:

- Specify the physical interface name.

[edit protocols [ancp](#)]

user@host# **set** [auto-configure-trigger interface](#) *physical-interface-name*

Release History Table

Release	Description
16.1R4	Starting in Junos OS Release 16.1R4, you can configure the ANCP agent to associate a neighbor with an access-facing physical interface for the creation of autosensed dynamic VLANs on the interface.

RELATED DOCUMENTATION

[Configuring the ANCP Agent](#)[Configuring ANCP Neighbors](#)[ANCP and the ANCP Agent Overview](#)[Layer 2 Wholesale with ANCP-Triggered VLANs Overview | 116](#)[Configuring Out-of-Band ANCP Messages to Trigger Dynamic VLAN Instantiation | 141](#)

Configuring a Username for Authentication of Out-of-Band Triggered Dynamic VLANs

When a subscriber logs in, the Access-Request message that is sent to the RADIUS server includes a username and optionally a password generated locally on the router to authenticate the subscriber during the VLAN authorization process. For a Layer 2 network that is wholesaled to a retailer where the dynamic VLANs are instantiated by out-of-band ANCP Port Up messages, you can configure the router to create a unique username with the value of the ANCP TLVs—Access-Loop-Circuit-ID, Access-Loop-Remote-Id, or both—as received in the ANCP Port Up message from the access node.

This configuration assumes the following:

- The ANCP agent is configured to notify AAA when it receives ANCP Port Up and Port Down messages.
- The dynamic profile is configured to instantiate a dynamic VLAN when notified by the ANCP agent that it has received an out-of-band ANCP Port Up message.
- The RADIUS authentication server is properly configured.

To include ANCP TLVs in the authentication username

1. (Optional) Specify inclusion of the Access-Loop-Circuit-ID TLV value.

```
[edit interfaces ge-0/0/0 auto-configure vlan-ranges username-include (Interfaces)]  
user@host# set circuit-id
```

2. (Optional) Specify inclusion of the Access-Loop-Remote-ID TLV value.

```
[edit interfaces ge-0/0/0 auto-configure vlan-ranges username-include (Interfaces)]  
user@host# set remote-id
```

NOTE: This ANCP information is not supported in stacked VLANs.

NOTE: You can use any of the attributes available to the **username-include** statement, except: **mac-address**, **option-18**, **option-37**, and **option-82**.

You can include other information in the username as for conventional autosensed dynamic VLANs. Alternatively, if you configure the router to convey ANCP-sourced access loop attributes as Juniper Networks VSAs—in this case Acc-Loop-Cir-Id (26-110) and Acc-Loop-Remote-Id (26-182)—the Access-Request message includes sufficient unique access line information for the RADIUS server to determine whether the access loop is wholesaled to a retailer or retained for the wholesaler.

RELATED DOCUMENTATION

Configuring VLAN Interface Username Information for AAA Authentication

[Configuring Out-of-Band ANCP Messages to Trigger Dynamic VLAN Instantiation | 141](#)

[Configuring the ANCP Agent for ANCP-Triggered, Autosensed Dynamic VLANs | 138](#)

[Layer 2 Wholesale with ANCP-Triggered VLANs Overview | 116](#)

Configuring Out-of-Band ANCP Messages to Trigger Dynamic VLAN Instantiation

The instantiation of conventional autosensed dynamic VLANs is triggered by in-band PPPoE or DHCP control packets that the Packet Forwarding Engine exceptions to the Routing Engine. A VLAN is authorized based on information extracted from specific fields and created according to a dynamic profile assigned to the VLAN range or stacked VLAN range.

Another way to instantiate an autosensed dynamic VLAN is with the processing of packets from an out-of-band protocol, ANCP. The out-of-band protocol method is useful where the traffic received might not be PPPoE or DHCP, such as in a Layer 2 wholesale scenario, where the traffic for an entire outer VLAN is wholesaled to a retailer and the VLANs are based on access line identifiers.

For this method, you configure the dynamic profile to accept packets from the out-of-band protocol. The dynamic profile is on an access-facing physical interface and is associated with a VLAN range available for the autosensed VLANs.

This configuration assumes the following:

- The dynamic profile is configured to instantiate a dynamic VLAN when notified by the ANCP agent that it has received an out-of-band ANCP Port Up message.
- The RADIUS authentication server is properly configured to authorize the VLANs and apply services as needed.
- The ANCP agent is configured to notify AAA when it receives ANCP Port Up and Port Down messages.
- The ANCP agent is configured to initiate interim accounting updates (which also enables immediate interim accounting updates) in response to information received in Port Up messages.

NOTE: Out-of-band triggering is supported only for single-tag VLANs; it is not supported for stacked VLANs.

To configure the instantiation of autosensed dynamic VLANs by out-of-band ANCP packets:

- Specify that ANCP packets are accepted.

```
[edit interfaces interface-name auto-configure vlan-ranges dynamic-profile profile-name]  
user@host# set accept-out-of-band ancp
```

RELATED DOCUMENTATION

[Layer 2 Wholesale with ANCP-Triggered VLANs Overview | 116](#)

[Configuring the ANCP Agent for ANCP-Triggered, Autosensed Dynamic VLANs | 138](#)

Configuring VLAN Interface Username Information for AAA Authentication

Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs

Triggering ANCP OAM to Simulate ANCP Port Down and Port Up Messages

You can trigger ANCP OAM to simulate the sending of an ANCP Port-Down or Port-Up message. Typically you use this feature only when troubleshooting an ANCP issue or to mitigate an error condition when ANCP is not operating normally.

When you issue either the **request ancp oam port-down** command or the **request ancp oam port-up** command from operational mode, you must specify either an IP address for an ANCP neighbor or the physical interface used for subscriber access. You must also specify all of the following; all three are required together to identify the access node:

- circuit-id *aci*—ANCP Access-Loop-Circuit-ID TLV
- remote-id *ari*—ANCP Access-Loop-Remote-ID TLV
- outer-vlan-id *vlan-id*—ANCP Access-Aggregation-Circuit-ID-Binary TLV

You can use the **request ancp oam port-up** command to trigger reauthorization and re-creation of the dynamic VLAN session and logical interface that is supporting Layer 2 wholesale after they have been removed by any of the following:

- Issuance of the **clear network-access aaa subscriber** command.
- Receipt of a RADIUS disconnect message that does not include the RADIUS Acct-Terminate-Cause attribute (49).
- Action by the ANCP agent.

The previous instance of the VLAN can be either ANCP-triggered (a wholesaled VLAN) or a conventionally autosensed dynamic VLAN (an access-provider-owned VLAN).

If no access line parameters are available from ANCP for a given access line, you can use the **request ancp oam port-up** command as a test mechanism to trigger authorization of a dynamic VLAN session and logical interface. The session and interface are created when a RADIUS Access-Accept message is subsequently received.

These commands have no effect on conventionally autosensed dynamic VLANs (for the access provider's own subscriber sessions) that have matching access loop attributes.

NOTE: Genuine ANCP Port-Down and Port-Up messages take precedence over these simulated messages. This means that when a Port-Down message has already been received, you cannot use the **request ancp oam port-up** command to initiate the Port-Up condition. When a Port-Up message has already been received, you cannot use the **request ancp oam port-down** command to initiate the Port-Down condition.

You can use the **request ancp oam port-down** command to trigger removal of the ANCP-triggered, autosensed, dynamic VLAN that corresponds to the specified attributes. The typical use for this command is to remove the VLAN created by sending a **request ancp oam port-up** command.

To simulate an ANCP Port Up message:

- Identify the loop by the neighbor's IP address or the access-facing physical interface, and the ACI, ARI, and outer VLAN ID.

```
user@host> request ancp oam port-up neighbor 192.168.32.5 circuit-id line-aci-1 remote-id line-ari-1
outer-vlan-id 126
user@host> request ancp oam port-up subscriber-interface ge-1/0/1 circuit-id line-aci-1 remote-id line-ari-1
outer-vlan-id 126
```

To simulate an ANCP Port Down message:

- Identify the loop by the neighbor's IP address or the access-facing physical interface, and the ACI, ARI, and outer VLAN ID.

```
user@host> request ancp oam port-down neighbor 192.168.32.5 circuit-id line-aci-1 remote-id line-ari-1
outer-vlan-id 126
user@host> request ancp oam port-down subscriber-interface ge-1/0/1 circuit-id line-aci-1 remote-id line-ari-1
outer-vlan-id 126
```

To verify the operation of either request, you can enter the following commands before and after initiating the Port Down or Port Up message:

- **show subscribers client-type vlan-oob detail**—Subscriber information is displayed for the VLAN on Port UP, or disappears on Port Down.
- **show subscribers summary**—The VLAN-OOB counter reflects the creation or removal of the VLAN-OOB session by incrementing (Port Up) or decrementing (Port Down).
- **show l2-routing-instance routing-instance-name**—The VLAN counters reflect to reflect the creation or removal of the VLAN-OOB session by incrementing (Port Up) or decrementing (Port Down).

RELATED DOCUMENTATION

[Layer 2 Wholesale with ANCP-Triggered VLANs Overview](#) | 116

[Triggering ANCP OAM to Test the Local Loop](#)

Configuring the ANCP Agent to Dampen the Effects of Short-Term Adjacency Losses

By default, the ANCP agent treats a loss of adjacency as if it has received a Port Down message for every access loop that is represented by the adjacency. All Layer 2 wholesale sessions are logged out and cleaned up. If the associated physical interface is in the Down state, then pending sessions cannot be reestablished when the interface transitions back to the Up state.

You can configure the ANCP agent to maintain the corresponding ANCP-triggered Layer 2 wholesale sessions for a configurable period in the event that an ANCP adjacency is lost. If the adjacency is restored before the timer expires, the session continues. If the timer expires before the adjacency is restored, then the session is logged out and cleaned up. This behavior dampens the effect of unstable ANCP connections. The hold timer can also detect when an access line is unconfigured on a neighbor and trigger logout and cleanup of the related sessions.

NOTE: The default value of the timer is 0, which means that the loss of neighbor adjacency immediately triggers a logout of all corresponding Layer 2 wholesale sessions.

To configure how long the ANCP agent maintains sessions in the event of an adjacency loss for any neighbor:

- Specify the hold timer duration in seconds.

```
[edit protocols ancp]
user@host# set adjacency-loss-hold-time seconds
```

To configure how long the ANCP agent maintains sessions in the event of an adjacency loss for a specific neighbor:

- Specify the hold timer duration in seconds.

```
[edit protocols ancp neighbor ip-address]
user@host# set adjacency-loss-hold-time seconds
```

RELATED DOCUMENTATION

[Configuring the ANCP Agent](#)[Configuring ANCP Neighbors](#)[ANCP and the ANCP Agent Overview](#)[Layer 2 Wholesale with ANCP-Triggered VLANs Overview | 116](#)

Reestablishing Pending Access Line Sessions for Layer 2 Wholesale

The access lines for ANCP-triggered, Layer 2 wholesale sessions can transition to a pending state after an ANCP adjacency loss when the inner VLAN ID swap range has been exhausted of tags and no other eligible core-facing physical interfaces are available. Typically, the sessions are reestablished when more VLAN IDs are made available, such as by extending the swap range, or more interfaces are available, such as by reconfiguration. When that does not happen, you can manually initiate the reestablishment process by issuing the **request auto-configuration reconnect-pending** command.

To manually reestablish sessions for which the corresponding access lines are in the pending state:

- Specify the routing instance with the reconnection request.

```
user@host> request auto-configuration reconnect-pending
```

RELATED DOCUMENTATION

[Layer 2 Wholesale with ANCP-Triggered VLANs Overview | 116](#)

Configuring Multiple Non-Overlapping VLAN Ranges for Core-Facing Physical Interfaces

You can configure up to 32 non-overlapping inner VLAN ID swap ranges for each core-facing physical interface in a Layer 2 wholesale network with VLAN-OOB subscribers. VLAN IDs from the ranges are allocated to replace the outer VLAN tag on traffic received on the access-facing physical interfaces. The swap occurs before the subscriber traffic is forwarded to the network service provider (NSP).

You can add or remove ranges or increase or decrease the size of existing ranges even while Layer 2 wholesale sessions are assigned to the core-facing interface associated with the ranges. You cannot remove a range from which a VLAN ID has already been allocated. You cannot reduce a range if the new range excludes a VLAN ID that has already been allocated.

To configure multiple ranges per interface:

- Specify the ranges.

```
user@host# set interfaces interface-name unit logical-unit-number inner-vlan-id-swap-ranges
    low-inner-tag1-high-inner-tag1
user@host# set interfaces interface-name unit logical-unit-number inner-vlan-id-swap-ranges
    low-inner-tag2-high-inner-tag2
user@host# set interfaces interface-name unit logical-unit-number inner-vlan-id-swap-ranges
    low-inner-tag3-high-inner-tag3
...
```

You can configure the ranges in any order. For example, one way to configure three non-overlapping ranges on interface ge-0/1/1 is the following:

```
[edit]
user@host# set interfaces ge-0/1/1 unit 0 inner-vlan-id-swap-ranges 70-80
user@host# set interfaces ge-0/1/1 unit 0 inner-vlan-id-swap-ranges 100-120
user@host# set interfaces ge-0/1/1 unit 0 inner-vlan-id-swap-ranges 10-60
```

Regardless of the order of configuration, **show** commands display the ranges in ascending order from lowest to highest:

```
user@host> show interfaces ge-0/1/1
```

```
description "ISP 1 core-facing PE1";
encapsulation ethernet-vpls;
unit 0 {
    inner-vlan-id-swap-ranges [10-60 70-80 100-120];
    ...
}
```

RELATED DOCUMENTATION

Clearing ANCP Access Loops

You can force a reset of a particular Layer 2 wholesale connection while the access loop is operationally up by issuing the **clear ancp access-loop** command. The command initiates logout of an ANCP-triggered, dynamic VLAN session, which includes issuing RADIUS Accounting-Stop messages for the session, and removal of the dynamic VLAN logical interface and active services. After the session is cleaned up, the command initiates re-authorization of the dynamic VLAN session, simulating receipt of an ANCP Port Up message. The session may then be recreated.

You must identify the access loop by either the IP address of the ANCP neighbor or the name of the subscriber-facing physical interface. You must also specify one or more of the following additional identifiers for the access loop:

- **circuit-id**—The value of the ANCP Access-Loop-Circuit-ID TLV.
- **remote-id**—The value of the ANCP Access-Loop-Remote-ID TLV.
- **outer-vlan-id**—The value of the ANCP Access-Aggregation-Circuit-ID-binary TLV.

NOTE:

The **clear ancp access-loop** command has no effect in the following circumstances:

- The access line is reported to be down, as indicated by an ANCP Port Down message, when the command is issued.
- An ANCP Port Down message is received for the access line while the dynamic VLAN logical interface and the services are being removed. In this case, re-authorization of the dynamic VLAN cannot take place until an ANCP Port Up message is received for that access line.
- A conventionally autosensed dynamic VLAN (for the access provider's own subscriber sessions) has matching access loop attributes. In this case, the Layer 2 wholesale access line for which the command is intended is cleared, but the other VLAN, for sessions owned by the access-provider, is cleared as expected.

To clear an ANCP access loop:

- Identify the loop by the neighbor's IP address or the access-facing physical interface, and one or more of the ACI, ARI, and outer VLAN ID.

```
user@host> clear ancp access-loop neighbor 192.168.32.5 circuit-id line-aci-1 remote-id line-ari-1 outer-vlan-id 126
user@host> clear ancp access-loop subscriber-interface ge-1/0/1 circuit-id line-aci-1 remote-id line-ari-1 outer-vlan-id 126
```

RELATED DOCUMENTATION

| [Layer 2 Wholesale with ANCP-Triggered VLANs Overview](#) | 116

Configuring Flat-File Accounting for Layer 2 Wholesale Services

IN THIS CHAPTER

- [Flat-File Accounting Overview | 149](#)
- [Configuring Flat-File Accounting for Layer 2 Wholesale | 153](#)
- [Configuring Flat-File Accounting for Extensible Subscriber Services Management | 158](#)
- [Configuring Service Accounting in Local Flat Files | 163](#)

Flat-File Accounting Overview

Accounting statistics can be collected from the Packet Forwarding Engine and reported in an XML flat file, which both contains and describes the data. Starting in Junos OS Release 16.1R4, you can use a flat-file profile that acts as a template to define attributes for accounting flat files.

Subscriber service accounting statistics are typically collected based on RADIUS Acct-Start and Acct-Stop messages that are sent to a RADIUS server individually or in bulk.

Starting in Junos OS Release 17.1R1, you can alternatively configure service-filter-based accounting statistics to be recorded per subscriber in a local flat file that is not automatically forwarded to a RADIUS server. This configuration collects the running total service statistics per interface family. Service accounting is initiated when the service profile is attached to the interface, whether by a static configuration or a RADIUS Change of Authorization (CoA) message.

NOTE: Starting in Junos OS Release 18.4R1, flat-file service accounting to a local file is no longer supported.

When the accounting file is created, a file header is also created if the file format is IP Detail Record (IPDR). The header is not created if the format is comma-separated variable (CSV). The file header includes the following information:

- XML namespace—Static link to the World Wide Web Consortium (W3C) organization's XML Schema Instance (XSI) definition.
- Schema version—Configurable name of the schema that defines the information conveyed in the accounting file. The schema version is associated with a specific XML format and output based on the flat-file profile configuration that is used for the business purpose. This structure enables the XML-formatted contents of the file to be correctly interpreted by the service provider's external file processor.
- NAS ID—Name of the BNG host (network access server) where the accounting statistics are collected.
- File creation timestamp—UTC time zone date and time when the accounting file was created.
- File ID—Number identifying the file. The ID is incremented when a new accounting file is created and can range from 1 through 2,147,483,647.

For example, consider the following sample header for an accounting file for Extensible Subscriber Services Manager (ESSM) business subscribers:

```
<BNGFile xmlns:xsi=http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="BNG_IPDR_20130423.xsd" NAS-ID="host-mx480-x5"
  FileCreationTimeStamp="2015-10-09T08:25:50" FileID="29">
<IPDR>
.....
.....
</IPDR>
</BNGFile>
```

Table 9 on page 150 lists the elements and their values in the sample header.

Table 9: Value of Elements in Sample Accounting Flat File XML Header

Description	Header Element	Value
XML namespace	xmlns	:xsi=http://www.w3.org/2001/XMLSchema-instance"
schema version	xsi:noNamespaceSchemaLocation	BNG_IPDR_20130423.xsd
NAS ID	NAS-ID	host-mx480-x5
file creation timestamp	FileCreationTimeStamp	2015-10-09T08:25:50
File ID	FileID	29

You can configure the following options for flat-file accounting at the **[edit accounting-options file filename]** hierarchy level:

- Maximum size of the accounting file.
- Number of files that are saved before overwriting.
- One or more sites where the files are sent for archiving.
- Frequency at which the files are transferred to an archive site.
- Start time for file transfer.
- Compression for the transferred files.
- Local backup on the router for files when transfer fails.
- Whether accounting files are saved when a change in mastership occurs for both the new master Routing Engine and the new backup Routing Engine or for only the new master Routing Engine.
- How long files are kept before being deleted from the local backup directory.

You can also create one or more flat-file profiles at the **[edit accounting-options flat-file-profile profile-name]** hierarchy level that act as templates to specify the following attributes for new accounting files when they are created:

- Statistics fields that you want to collect, such as egress statistics or ingress statistics fields.
- Name and format of the accounting file.
- Frequency at which the Packet Forwarding Engine is polled for the statistics.
- Schema version.

Archive sites provide security and storage for the accounting files, which are transferred at regular intervals. When more than one archive site is configured, the router attempts to transfer the files to the first site on the list. If that fails, the router tries each of the other sites in turn until the transfer either succeeds for one site or fails for all sites. If you configure the last site in the list to be a local directory on the router rather than another remote site, then the files are backed up locally if all remote sites fail. The failed files are simply stored in the designated site. They are not automatically resubmitted to the archival sites. You must use an event script or some other means to have these files resubmitted. Any files remaining in the local directory are deleted when the **cleanup-interval** expires.

Alternatively, you can use the **backup-on-failure** statement at the **[edit accounting-options file filename]** hierarchy level to back up the files locally if all the remote attempts fail. If that occurs, the router compresses the accounting files and backs them up to the **/var/log/pfedBackup/** directory. Whenever any of the archive sites is reachable, the router attempts to transfer the data from **/var/log/pfedBackup/** to that site in compressed format. If the transfer of the backed-up files to the reachable site fails, the system tries to transfer the files to any other site that becomes reachable during the transfer interval. Any files that fail to transfer are compressed and kept in **/var/log/pfedBackup/** until an archival site is reachable and the

files are successfully transferred. Any files that remain in that directory are deleted when the **cleanup-interval** expires.

BEST PRACTICE: Use the **backup-on-failure** feature to reliably and automatically back up files and retransmit them to archives rather than relying on a local site listed as the last archive site.

If the backup Routing Engine does not have access to the archive site—for example, when the site is not connected by means of an out-of-band interface or when the path to the site is routed through a line card—you can ensure that the backup Routing Engine's accounting files are backed up by using the **push-backup-to-master** statement at the **[edit accounting-options file filename]** hierarchy level. When a change in mastership occurs, the new backup Routing Engine saves its files to the **/var/log/pfedBackup/** directory. The master Routing Engine subsequently includes these files when it sends its own accounting files to the archive site at every transfer interval.

To conserve resources during transfer of accounting files and at the archive site, use the **compress** statement at the **[edit accounting-options file filename]** hierarchy level to compress the files when they are transferred. This option is disabled by default.

A system logging message is generated when a transfer succeeds (**transfer-file: Transferred filename**) or fails (**transfer-file failed to transfer**). In the event of a failure, an error message is logged to indicate the nature of the failure.

Release History Table

Release	Description
17.1R1	Starting in Junos OS Release 17.1R1, you can alternatively configure service-filter-based accounting statistics to be recorded per subscriber in a local flat file that is not automatically forwarded to a RADIUS server.
16.1R4	Starting in Junos OS Release 16.1R4, you can use a flat-file profile that acts as a template to define attributes for accounting flat files.

RELATED DOCUMENTATION

[Configuring Flat-File Accounting for Extensible Subscriber Services Management](#) | 158

[Configuring Flat-File Accounting for Layer 2 Wholesale](#) | 153

[Configuring Service Accounting in Local Flat Files](#) | 163

[Configuring Accounting-Data Log Files](#)

Configuring Flat-File Accounting for Layer 2 Wholesale

Flat-file accounting is typically used for recording accounting statistics on logical interfaces for Extensible Subscriber Services Manager (ESSM) business subscribers. However, starting in Junos OS Release 16.1R4, you can also use flat-file accounting to collect and archive various accounting statistics for your Layer 2 wholesale environment. You do this by creating a flat-file profile and applying it to a core-facing physical interface.

You can also configure a flat-file profile to monitor and report Layer 2 multicast statistics; you assign this profile to the logical interface configured on the core-facing physical interface. This approach enables you to have separate accounting files that overlap in content only in the non-statistical, general parameters. The Layer 2 multicast statistics are available only when the encapsulation on the logical interface is **ethernet-vpls**.

You can configure multiple accounting profiles with different combinations of fields for specific accounting requirements, and then assign the profiles as needed to provisioned interfaces to satisfy the accounting requirements for each interface depending on how the interface is used.

A given flat-file profile can be assigned to both use cases; for example, by specifying **all-fields** for a global or group level. In this case, the fields you configure appear in the accounting record only if they make sense in the context.

BEST PRACTICE: We recommend you use separate flat-file profiles for Layer 2 wholesale core-facing physical interfaces and ESSM business subscriber logical interfaces.

Some statistics and general parameter fields are available either only for logical interfaces or only for physical interfaces. The **accounting-type**, **line-id**, **nas-port-id**, and **vlan-id** general parameters are not available for core-facing physical interfaces. Because the core-facing physical interfaces carry Layer 2 cross-connected sessions, no useful IPv6 statistics are available. Accordingly, do not configure the **input-v6-bytes**, **input-v6-packets**, **output-v6-bytes**, or **output-v6-packets** overall packet fields.

To configure flat-file accounting for a Layer 2 wholesale network:

1. Create a flat-file profile.

```
[edit accounting-options]
user@host# edit flat-file-profile profile-name
```

2. (Optional) Configure the name of the XML schema for the accounting flat file.

```
[edit accounting-options flat-file-profile profile-name]
```

```
user@host# set schema-version schema-name
```

3. Specify the filename for the accounting file.

```
[edit accounting-options flat-file-profile profile-name]  
user@host# set file accounting-filename
```

4. Specify the general, nonstatistical parameters for the accounting file that are displayed as part of the accounting record header.

```
[edit accounting-options flat-file-profile profile-name fields]  
user@host# set general-param option
```

BEST PRACTICE: We recommend that you include the general parameter **all-fields** option for both core-facing physical interfaces and, when you are collecting Layer 2 multicast statistics, on the logical interface that represents the physical interface.

5. Specify the accounting statistics that are collected and recorded in the accounting file for the core-facing physical interface.

```
[edit accounting-options flat-file-profile profile-name fields]  
user@host# set egress-stats option  
user@host# set ingress-stats option  
user@host# set overall-packet option
```

BEST PRACTICE: We recommend that you include the following statistics fields in flat-file profiles for core-facing physical interfaces:

- Egress statistics fields: **all-fields**
- Ingress statistics fields: **all-fields**
- Overall packet fields: **input-bytes, input-discards, input-errors, input-packets, output-bytes, output-errors, output-packets**

6. (Optional) For Layer 2 multicast statistics, specify the accounting statistics that are collected and recorded in the accounting file for the logical interface representing the core-facing physical interface.

```
[edit accounting-options flat-file-profile profile-name fields]
user@host# set l2-stats option
```

BEST PRACTICE: We recommend that you include the following statistics fields in flat-file profiles for logical interfaces on the core-facing physical interfaces:

- Layer 2 statistics fields: **all-fields**

7. (Optional) Specify the format of the accounting file.

```
[edit accounting-options flat-file-profile profile-name]
user@host# set format (csv | ipdr)
```

8. (Optional) Specify the interval at which the Packet Forwarding Engine associated with the interface is polled for the statistics specified in the profile.

```
[edit accounting-options flat-file-profile profile-name]
user@host# set interval minutes
```

NOTE: When you do not configure this option, the polling interval is 15 minutes.

9. Configure the maximum size of the accounting file.

```
[edit accounting-options file filename]
user@host# set size bytes
```

10. Configure one or more archive sites for the files.

```
[edit accounting-options file filename]
user@host# set archive-sites site-name
```

The **site-name** is any valid FTP or Secure FTP URL. When the file is archived, the router attempts to transfer the file to the archive site. If you have specified more than one site, it tries the first site in the list. If that fails, it tries each site in turn until a transfer succeeds. The log file is stored at the archive site with a filename of the format **router-name_log-filename_timestamp**. The last site in a list is often a local directory, in case no remote site is reachable.

11. (Optional) Configure the start time for transferring files.

```
[edit accounting-options file filename]  
user@host# set start-time YYYY-MM-DD.hh:mm
```

12. (Optional) Configure how frequently the file is transferred.

```
[edit accounting-options file filename]  
user@host# set transfer-interval minutes
```

NOTE: When you do not configure this option, the file is transferred every 30 minutes.

13. (Optional) Configure the maximum number of files (3 through 1000) to save.

```
[edit accounting-options file filename]  
user@host# set files number
```

NOTE: When you do not configure this option, a maximum of 10 files are saved.

14. (Optional) Configure the router to save a backup copy of the accounting file to the `/var/log/pfedBackup` directory if the normal transfer of the files to the archive sites fails.

```
[edit accounting-options file filename]  
user@host# set backup-on-failure
```

NOTE: When you do not configure this option, the file is saved on failure into the local directory specified as the last site in the list of archive sites.

15. (Optional) Configure the accounting file to be compressed during transfer to an archive site.

```
[edit accounting-options file filename]  
user@host# set compress
```


16. (Optional) Configure the router's new backup Routing Engine to send its accounting file to the **/var/log/pfedBackup** directory on the new master Routing Engine when a change in mastership occurs.

```
[edit accounting-options file filename]
user@host# set push-backup-to-master
```

17. (Optional) Configure the number of days after which accounting files are deleted from the local backup directory.

```
[edit accounting-options]
user@host# set cleanup-interval days
```

NOTE: Files are retained for 1 day if you do not configure this option.

18. Assign the profile to the relevant interface.

For the core-facing physical interface:

```
[edit interfaces physical-interface-name]
user@host# set accounting-profile flat-file-profile-name
```

For the logical interface representing the core-facing physical interface:

```
[edit interfaces physical-interface-name unit logical-unit-number]
user@host# set accounting-profile flat-file-profile-name
```

Release History Table

Release	Description
16.1R4	However, starting in Junos OS Release 16.1R4, you can also use flat-file accounting to collect and archive various accounting statistics for your Layer 2 wholesale environment.

RELATED DOCUMENTATION

[Configuring Accounting-Data Log Files](#)

[Flat-File Accounting Overview](#) | 149

Configuring Flat-File Accounting for Extensible Subscriber Services Management

Flat-file accounting is typically used to collect and archive various accounting statistics on logical interfaces for Extensible Subscriber Services Manager (ESSM) business subscribers. Other applications include accounting for wholesaler and retailer subscriber activity in a Layer 2 wholesale environment. Starting in Junos OS Release 16.1R4, you can create a flat-file profile to use as a template to define attributes for accounting flat files. The profile specifies the following:

- The statistics fields that are collected.
- The filename where the statistics are logged.
- The format of the file, the interval at which the statistics are collected.
- The name of the XML schema file that specifies the contents of the accounting file.

You can configure multiple accounting profiles with different combinations of fields for specific accounting requirements, and then assign the profiles as needed to provisioned interfaces to satisfy the accounting requirements for each interface depending on how it is used.

A given flat-file profile can be assigned to both use cases; for example, by specifying **all-fields** for a global or group level. In this case, the fields you configure appear in the accounting record only if they make sense in the context.

BEST PRACTICE: We recommend you use separate flat-file profiles for ESSM business subscriber logical interfaces and Layer 2 wholesale core-facing physical interfaces.

To configure flat-file accounting for ESSM business services:

1. Create a flat-file profile.

```
[edit accounting-options]  
user@host# edit flat-file-profile profile-name
```

2. (Optional) Configure the name of the XML schema for the accounting flat file.

```
[edit accounting-options flat-file-profile profile-name]
user@host# set schema-version schema-name
```

3. Specify the filename for the accounting file.

```
[edit accounting-options flat-file-profile profile-name]
user@host# set file accounting-filename
```

4. Specify the general, nonstatistical parameters for the accounting file that are displayed as part of the accounting record header.

```
[edit accounting-options flat-file-profile profile-name fields]
user@host# set general-param option
```

BEST PRACTICE: We recommend that you include the following general parameter fields in flat-file profiles for ESSM subscribers:

- General parameter fields: **accounting-type**, **descr**, **line-id**, **logical-interface**, **nas-port-id**, **timestamp**, and **vlan-id**

5. Specify the accounting statistics that are collected and recorded in the accounting file.

```
[edit accounting-options flat-file-profile profile-name fields]
user@host# set egress-stats option
user@host# set ingress-stats option;
user@host# set overall-packet option;
```

BEST PRACTICE: We recommend that you include the following statistics fields in flat-file profiles for core-facing physical interfaces:

- Egress statistics fields: **all-fields**
- Ingress statistics fields: **all-fields**
- Overall packet fields: **all-fields**

6. (Optional) Specify the format of the accounting file.

```
[edit accounting-options flat-file-profile profile-name]
user@host# set format (csv | ipdr)
```

7. (Optional) Specify the interval at which the Packet Forwarding Engine associated with the interface is polled for the statistics specified in the profile.

```
[edit accounting-options flat-file-profile profile-name]
user@host# set interval minutes
```

NOTE: When you do not configure this option, the polling interval is 15 minutes.

8. Configure the maximum size of the accounting file.

```
[edit accounting-options file filename]
user@host# set size bytes
```

9. Configure one or more archive sites for the files.

```
[edit accounting-options file filename]
user@host# set archive-sites site-name
```

The **site-name** is any valid FTP or Secure FTP URL. When the file is archived, the router attempts to transfer the file to the archive site. If you have specified more than one site, it tries the first site in the list. If that fails, it tries each site in turn until a transfer succeeds. The log file is stored at the archive site with a filename of the format **router-name_log-filename_timestamp**. The last site in a list is often a local directory, in case no remote site is reachable.

10. (Optional) Configure the start time for transferring the file.

```
[edit accounting-options file filename]
user@host# set start-time YYYY-MM-DD.hh:mm
```

11. (Optional) Configure how frequently the file is transferred.

```
[edit accounting-options file filename]
user@host# set transfer-interval minutes
```

NOTE: When you do not configure this option, the file is transferred every 30 minutes.

12. (Optional) Configure the maximum number of files (3 through 1000) to save.

```
[edit accounting-options file filename]
user@host# set files number
```

NOTE: When you do not configure this option, a maximum of 10 files are saved.

13. (Optional) Configure the router to save a backup copy of the accounting file to the `/var/log/pfedBackup` directory if the normal transfer of the files to the archive sites fails. Specify whether only the current file from the master Routing Engine is saved or both that file and the file from the backup Routing Engine.

```
[edit accounting-options file filename]
user@host# set backup-on-failure (master-and-slave | master-only)
```

NOTE: When you do not configure this option, the file is saved on failure into the local directory specified as the last site in the list of archive sites.

14. (Optional) Configure the accounting file to be compressed during transfer to an archive site.

```
[edit accounting-options file filename]
user@host# set compress
```

15. (Optional) Configure the router's new backup Routing Engine to send its accounting file to the `/var/log/pfedBackup` directory on the new master Routing Engine when a change in mastership occurs.

```
[edit accounting-options file filename]
user@host# set push-backup-to-master
```

16. (Optional) Configure the number of days after which accounting files are deleted from the local backup directory.

```
[edit accounting-options]
user@host# set cleanup-interval days
```

NOTE: Files are retained for 1 day if you do not configure this option.

17. Assign the profile to an ESSM subscriber.

```
[edit system services extensible-subscriber-services]
user@host# set flat-file-profile flat-file-profile-name
```

Release History Table

Release	Description
16.1R4	Starting in Junos OS Release 16.1R4, you can create a flat-file profile to use as a template to define attributes for accounting flat files.

RELATED DOCUMENTATION

- Configuring Accounting-Data Log Files
- Configuring Flat-File Accounting for Layer 2 Wholesale | 153

Configuring Service Accounting in Local Flat Files

Starting in Junos OS Release 17.1R1, you can configure flat-file accounting to collect service statistics for subscribers and report those statistics to a local file. This configuration collects the running total service statistics per interface family. Because the statistics are maintained in the Routing Engine in a statistics database, they are not affected by a line-card restart, a graceful Routing Engine switchover, or a unified in-service software upgrade (ISSU). The statistics counters are reset when the router reboots.

NOTE: Starting in Junos OS Release 18.4R1, service accounting in local flat files is no longer supported. If included in a configuration, it is ignored.

To configure local flat-file accounting for services:

1. Configure the subscriber access profile to report service accounting records in a local flat file.

```
[edit access profile profile-name]  
user@host# set service accounting-order local
```

NOTE: When you configure **local**, the CLI checks at commit that the flat-file profile is configured under **[edit access profile *profile-name* local]**.

Alternatively, you can set the service accounting order to **activation-protocol** instead of **local**:

```
user@host# set service accounting-order activation-protocol
```

Use this option only when you plan to activate the service by means of the CLI configuration or a command. In this case, the CLI does not check for the flat-file profile to be configured. If the profile is not configured, no statistics are collected.

NOTE: When you configure the **local** option, both volume and time statistics are collected for the service accounting sessions. In this case, you must not configure the **volume-time** option at the **[edit access profile *profile-name* service accounting statistics]** hierarchy level; otherwise, an error is generated when you commit the configuration.

2. Specify the name of the flat-file profile that is used to collect the service statistics.

```
[edit access profile profile-name]
user@host# set local flat-file-profile flat-file-profile-name
```

3. Create the flat-file profile to collect the subscriber service accounting statistics and other parameters.

```
[edit accounting-options]
user@host# edit flat-file-profile profile-name
```

4. Specify the filename for the accounting file.

```
[edit accounting-options flat-file-profile profile-name]
user@host# set file accounting-filename
```

5. Specify that service accounting statistics are collected.

```
[edit accounting-options flat-file-profile profile-name fields]
user@host# set service-accounting
```

6. (Optional) Specify the general, nonstatistical parameters for the accounting file that are displayed as part of the accounting record header.

```
[edit accounting-options flat-file-profile profile-name fields]
user@host# set general-param option
```

7. (Optional) Configure the name of the XML schema for the accounting flat file.

```
[edit accounting-options flat-file-profile profile-name]
user@host# set schema-version schema-name
```

8. (Optional) Specify the format of the accounting file.

```
[edit accounting-options flat-file-profile profile-name]
user@host# set format (csv | ipdr)
```

NOTE: When you do not configure this option, the format is **ipdr**.

9. (Optional) Specify the interval at which the Packet Forwarding Engine associated with the interface is polled for the statistics specified in the profile.

```
[edit accounting-options flat-file-profile profile-name]
user@host# set interval minutes
```

NOTE: When you do not configure this option, the polling interval is 15 minutes.

NOTE:

The interval value configured in the flat-file profile can be overridden by other interval values:

- The service accounting update interval configured at the **edit access profile profile-name service accounting update-interval**] hierarchy level.
- An update interval value configured in the RADIUS attribute, Service-Interim-Acct-Interval (VSA 26–140). This value also overrides the service accounting update interval.

10. Configure the maximum size of the accounting file.

```
[edit accounting-options file filename]
user@host# set size bytes
```

11. (Optional) Configure the maximum number of files (3 through 1000) to save.

```
[edit accounting-options file filename]
user@host# set files number
```

NOTE: When you do not configure this option, a maximum of 10 files are saved.

12. (Optional) Configure one or more archive sites for the files.

```
[edit accounting-options file filename]
user@host# set archive-sites site-name
```

The **site-name** is any valid FTP or Secure FTP URL. When the file is archived, the router attempts to transfer the file to the archive site. If you have specified more than one site, it tries the first site in the

list. If that fails, it tries each site in turn until a transfer succeeds. The log file is stored at the archive site with a filename of the format **router-name_log-filename_timestamp**. The last site in a list is often a local directory, in case no remote site is reachable.

13. (Optional) Configure the start time for transferring files.

```
[edit accounting-options file filename]
user@host# set start-time YYYY-MM-DD.hh:mm
```

14. (Optional) Configure how frequently the file is transferred.

```
[edit accounting-options file filename]
user@host# set transfer-interval minutes
```

NOTE: When you do not configure this option, the file is transferred every 30 minutes.

15. (Optional) Configure the router to save a backup copy of the accounting file to the **/var/log/pfedBackup** directory if the normal transfer of the files to the archive sites fails.

```
[edit accounting-options file filename]
user@host# set backup-on-failure
```

NOTE: When you do not configure this option, the file is saved on failure into the local directory specified as the last site in the list of archive sites.

16. (Optional) Configure the accounting file to be compressed during transfer to an archive site.

```
[edit accounting-options file filename]
user@host# set compress
```

17. (Optional) Configure the router's new backup Routing Engine to send its accounting file to the **/var/log/pfedBackup** directory on the new master Routing Engine when a change in mastership occurs.

```
[edit accounting-options file filename]
user@host# set push-backup-to-master
```

18. (Optional) Configure the number of days after which accounting files are deleted from the local backup directory.

```
[edit accounting-options]
user@host# set cleanup-interval days
```

NOTE: When you do not configure this option, files are retained for only 1 day.

Release History Table

Release	Description
17.1R1	Starting in Junos OS Release 17.1R1, you can configure flat-file accounting to collect service statistics for subscribers and report those statistics to a local file.

RELATED DOCUMENTATION

Configuring Accounting-Data Log Files
[Flat-File Accounting Overview](#) | 149

Configuring Five-Level and Four-Level Heterogeneous Networks

IN THIS CHAPTER

- [Five-Level and Four-Level Heterogeneous Networks | 168](#)
- [OLT Migration to Using PON TLVs Instead of DSL TLVs | 190](#)

Five-Level and Four-Level Heterogeneous Networks

IN THIS SECTION

- [CoS Node Shaping in Four-Level and Five-Level Heterogeneous Networks | 168](#)
- [CuTTB Use Case Topology and CoS Hierarchy | 173](#)
- [FTTB/FTTH Use Case Topology and CoS Hierarchy | 178](#)
- [Automatic Creation of Business Subscriber Interface Sets | 183](#)
- [How to Configure the Automatic Creation of Business Subscriber Interface Sets | 184](#)
- [Dynamic Level 2 and Level 3 Interface Set Naming with Predefined Variables | 185](#)

CoS Node Shaping in Four-Level and Five-Level Heterogeneous Networks


A heterogeneous subscriber access model has the following characteristics:

- It includes both residential subscribers and business subscribers. Both subscriber types are typically PPPoE subscribers.
- The access technologies can be conventional or shared media, or both. Shared media access includes bonded copper connections through a DPU-C or fiber connections through a DPU-P. DPU-C and DPU-P are distribution units for the respective media type. Conventional access networks do not include either a DPU-C or DPU-P.

- Traffic shaping depends on hierarchical CoS. The network can use a four-level scheduler hierarchy, a five-level scheduler hierarchy, or both.

Figure 10 on page 169 summarizes how CoS shapes key nodes in the five-level scheduler hierarchy. Shaping is based either on the adjusted rates from the DSL and PON TLVs or on traffic control profiles in the dynamic client profile configuration. A CoS adjustment control profile specifies the source of the shaping rate applied to a given node.

Figure 10: Five-Level CoS Node Shaping Summary



The following lists describe the CoS scheduler nodes by access type and subscriber type for the five-level hierarchy in Figure 10 on page 169.

For conventional access, residential subscribers:

- Level 1 node—Corresponds to the access-facing physical interface.
- Level 2 node—Corresponds to a dynamic interface set that conserves L2 nodes. This parent interface set is based on the underlying physical interface. The name is derived from the predefined variable, \$junos-phy-ifd-underlying-intf-set-name, by appending “-underlying”. Traffic shaping is determined by a traffic control profile specified in the dynamic client profile.
- Level 3 node—Corresponds to a dynamic interface set that conserves Level 3 nodes. This child interface set is based on the physical interface. The name is derived from the predefined variable,

\$junos-phy-ifd-intf-set-name. Traffic shaping is determined by a traffic control profile specified in the dynamic client profile.

- Level 4 node—Corresponds to the subscriber's PPPoE session logical interface. Traffic shaping is determined by the Actual-Data-Rate-Downstream TLV (0x82).
- Level 5 node—Corresponds to the scheduling queue for the subscriber.

For conventional access, business subscribers:

- Level 1 node—Corresponds to the access-facing physical interface.
- Level 2 node—Corresponds to a dynamic interface set that conserves L2 nodes. This parent interface set is based on the underlying physical interface. The name is derived from the predefined variable, \$junos-phy-ifd-underlying-intf-set-name by appending "-underlying". Traffic shaping is determined by a traffic control profile specified in the dynamic client profile.
- Level 3 node—Corresponds to a dynamic interface set that conserves L3 nodes. This child interface set is based on the physical interface and VLAN tag. The set name is derived in one of two ways:
 - If configured, it is provided by the Qos-Set-Name VSA (26–4874–130) in the Access-Accept from the RADIUS server.
 - It is created from the \$junos-phy-ifd-interface-set-name predefined variable by appending the SVLAN tag to the value.

Traffic shaping is determined by the Actual-Data-Rate-Downstream TLV (0x82).

- Level 4 node—Corresponds to the subscriber's dynamic PPPoE session logical interface or static VLAN logical interface.
- Level 5 node—Corresponds to the scheduling queue for the subscriber.

For shared-media access, residential subscribers:

- Level 1 node—Corresponds to the access-facing physical interface.
- Level 2 node—Corresponds to a dynamic aggregation interface set that conserves L2 nodes. This parent interface set is based on the backhaul identifier from Access-Aggregation-Circuit-Id-ASCII TLV 0x03, which represents the PON tree connection. The name is derived from the predefined-variable, \$junos-aggregation-interface-set-name. Traffic shaping is determined by the Actual-Data-Rate-Downstream TLV (0x82) for bonded copper connections and by the PON-Tree-Maximum-Data-Rate-Downstream TLV (0x98) for PON tree connections.
- Level 3 node—Corresponds to a dynamic aggregation interface set that conserves L3 nodes. This child interface set is based on the backhaul identifier from the Access-Aggregation-Circuit-Id-ASCII TLV (0x03), which represents the PON tree connection. The name is derived from the predefined variable, \$junos-aggregation-interface-set-name by appending "-default". Traffic shaping is determined by a traffic control profile specified in the dynamic client profile.

- Level 4 node—Corresponds to the subscriber's PPPoE session logical interface. Traffic shaping is determined by the Actual-Data-Rate-Downstream TLV (0x82) for bonded copper connections and by the ONT/ONU-Peak-Data-Rate-Downstream TLV (0x94) for PON tree connections.
- Level 5 node—Corresponds to the scheduling queue for the subscriber.

For shared-media access, business subscribers:


- Level 1 node—Corresponds to the access-facing physical interface.
- Level 2 node—Corresponds to a dynamic aggregation interface set that conserves L2 nodes. This parent interface set is based on the backhaul identifier from the Access-Aggregation-Circuit-Id-ASCII TLV (0x03), which represents the PON tree connection. The name is derived from the predefined-variable, \$junos-aggregation-interface-set-name. Traffic shaping is determined by the Actual-Data-Rate-Downstream TLV (0x82) for bonded copper connections and by the PON-Tree-Maximum-Data-Rate-Downstream TLV (0x98) for PON tree connections.
- Level 3 node—Corresponds to a dynamic interface set that conserves L3 nodes. This child interface set is based on the physical interface and VLAN tag. The set name is derived in one of two ways:
 - If configured, it is provided by the Qos-Set-Name VSA (26-4874-130) in the Access-Accept from the RADIUS server.
 - It is created from the \$junos-phy-ifd-interface-set-name predefined variable by appending the SVLAN tag to the value.

Traffic shaping is determined by the Actual-Data-Rate-Downstream TLV (0x82) for bonded copper connections and by the ONT/ONU-Peak-Data-Rate-Downstream TLV (0x94) for PON tree connections.

- Level 4 node—Corresponds to the subscriber's dynamic PPPoE session logical interface or static VLAN logical interface.
- Level 5 node—Corresponds to the scheduling queue for the subscriber.

[Figure 11 on page 172](#) summarizes how CoS shapes key nodes in the four-level scheduler hierarchy. Shaping is based either on the adjusted rates resulting from the DSL and PON TLVs or on traffic control profiles in the dynamic client profile configuration. A CoS adjustment control profile specifies the source of the shaping rate applied to a given node.

Figure 11: Four-Level CoS Node Shaping Summary



The following lists describe the CoS scheduler nodes by access type and subscriber type for the four-level hierarchy in [Figure 11 on page 172](#).

For conventional access, residential subscribers:

- Level 1 node—Corresponds to the access-facing physical interface.
- Level 2 node—Corresponds to a dynamic interface set that conserves Level 2 nodes. This interface set is based on the physical interface. The name is derived from the predefined variable, `$junos-phy-ifd-intf-set-name`. Traffic shaping is determined by a traffic control profile specified in the dynamic client profile.
- Level 3 node—Corresponds to the subscriber's PPPoE session logical interface. Traffic shaping is determined by the Actual-Data-Rate-Downstream TLV (0x82).
- Level 4 node—Corresponds to the scheduling queue for the subscriber.

For conventional access, business subscribers:

- Level 1 node—Corresponds to the access-facing physical interface.
- Level 2 node—Corresponds to a dynamic interface set that conserves L2 nodes. This interface set is based on the physical interface and VLAN tag. The set name is derived in one of two ways:
 - If configured, it is provided by the Qos-Set-Name VSA (26–4874–130) in the Access-Accept from the RADIUS server.
 - It is created from the `$junos-phy-ifd-interface-set-name` predefined variable by appending the SVLAN tag to the value.

Traffic shaping is determined by the Actual-Data-Rate-Downstream TLV (0x82).

- Level 3 node—Corresponds to the subscriber's dynamic PPPoE session logical interface or static VLAN logical interface.
- Level 4 node—Corresponds to the scheduling queue for the subscriber.

For shared-media access, residential subscribers:

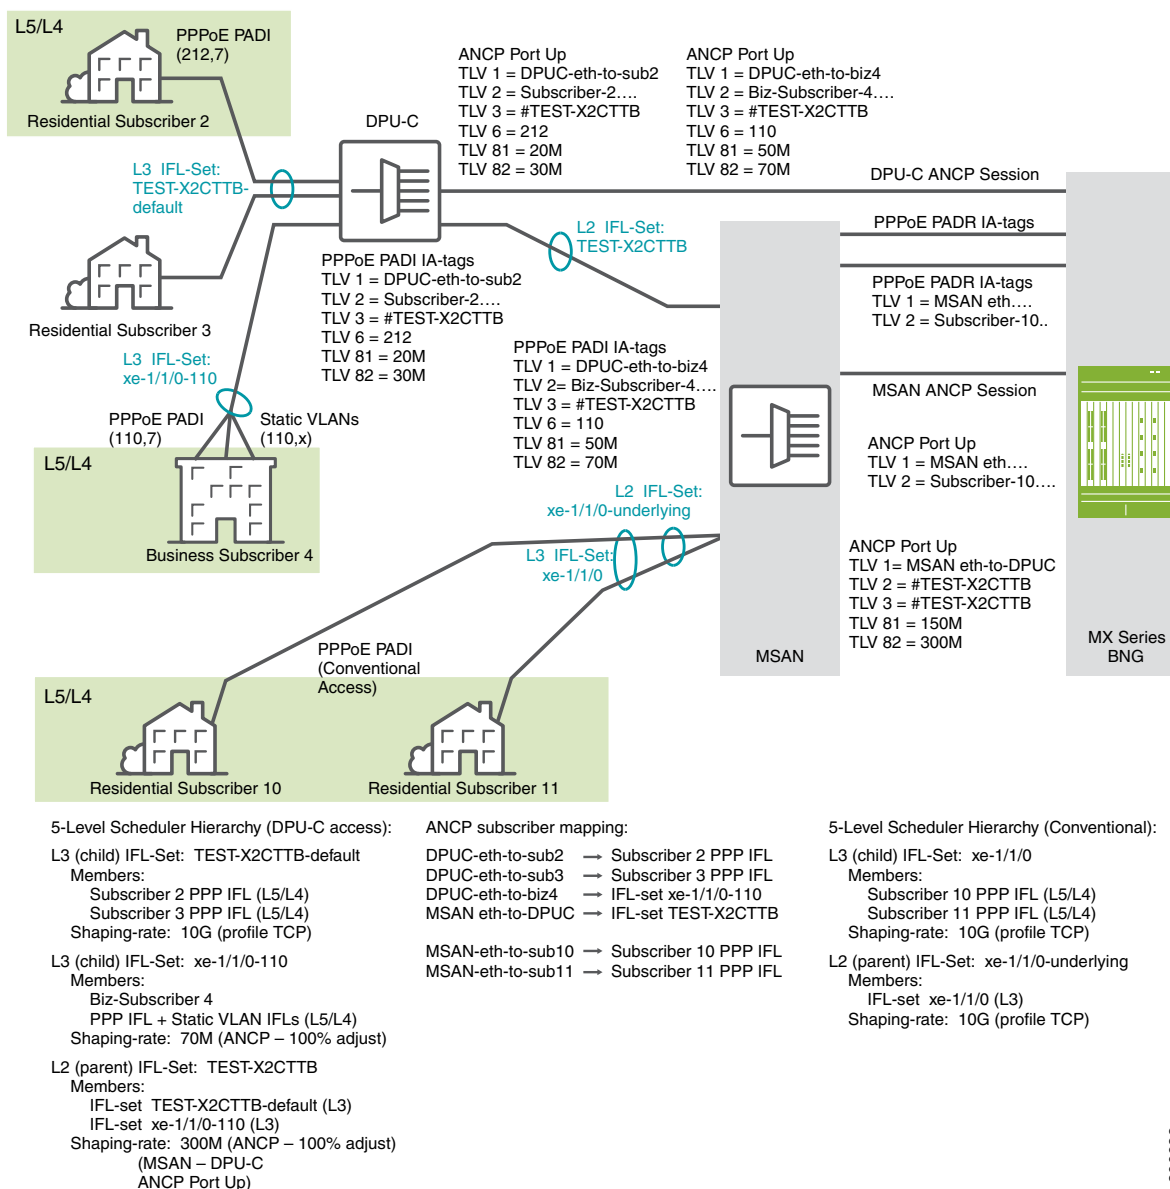
- Level 1 node—Corresponds to the access-facing physical interface.
- Level 2 node—Corresponds to a dynamic aggregation interface set that conserves Level 2 nodes. This interface set is based on the backhaul identifier from the Access-Aggregation-Circuit-Id-ASCII TLV (0x03), which represents the PON tree connection. The name is derived from the predefined-variable, \$junos-aggregation-interface-set-name. Traffic shaping is determined by the Actual-Data-Rate-Downstream TLV (0x82) for bonded copper connections and by the PON-Tree-Maximum-Data-Rate-Downstream TLV (0x98) for PON tree connections.
- Level 3 node—Corresponds to the subscriber's PPPoE session logical interface. Traffic shaping is determined by the Actual-Data-Rate-Downstream TLV (0x82) for bonded copper connections and by the ONT/ONU-Peak-Data-Rate-Downstream TLV (0x94) for PON tree connections.
- Level 4 node—Corresponds to the scheduling queue for the subscriber.

Business subscribers are not supported in a four-level, shared-media access network.

CuTTB Use Case Topology and CoS Hierarchy

[Figure 12 on page 174](#) shows a heterogeneous CuTTB topology that includes both shared-media (bonded copper through a DPU-C) and conventional, (nonbonded copper) access for PPPoE subscribers.

Figure 12: CuTTB CoS Hierarchy Example



This topology has the following subscribers:

- Two residential subscribers, 2 and 3, and a business subscriber, 4, have a shared-media access to the network through a DPU-C to the MSAN and then to the BNG.
- Two residential subscribers, 10 and 11, have conventional access to the network through an MSAN to the BNG.
- Residential subscriber 3 is not currently logged in.
- When residential subscriber 2 and business subscriber 4 log in:


1. PPPoE sends a PADI message to the DPU-C that includes the outer VLAN tag for each.
2. The DPU-C sends an ANCP Port Up message to the BNG for each subscriber. ANCP TLVs in the message identify the access line, the subscriber, the ASCII identifier for the access line, the VLAN outer tag as the binary identifier for the access line, the upstream rate, and the downstream rate.
The ASCII identifier (TLV 0x03) begins with the # character, signifying that the remainder of the value identifies the backhaul (bonded copper, shared media) line. TLV 0x03 is the same for both subscribers, because they connect through the same DPU-C.
3. The DPU-C sends a PADI message for each subscriber to the MSAN. The PADI conveys the PPPoE-IA tags that identify the same attributes as the ANCP TLVs.
4. The MSAN sends a PADR message with the PPPoE tags to the BNG. The MSAN also opens an ANCP session to the BNG by sending an ANCP Port Up message for the MSAN-to-DPU-C connection. The rates in TLVs 0x81 and 0x82 are the values for the MSAN-to-DPU-C line, represented by the L2 interface set. In other words, these are the rates for the bonded copper line itself rather than the subscriber access lines. TLV 0x03 value is also reported in TLV 0x02 to indicate the bonded copper line.

- When residential subscribers 10 and 11 log in:

1. PPPoE sends a PADI message for each subscriber to the MSAN. The PADI conveys the PPPoE-IA tags for the individual subscriber access lines.
2. The MSAN sends a PADR message with the PPPoE tags for each subscriber line to the BNG.
3. The MSAN also sends an ANCP Port Up message to the BNG for each subscriber. ANCP TLVs in the message identify the access line, the subscriber, the ASCII identifier for the access line, the VLAN outer tag as the binary identifier for the access line, the upstream rate, and the downstream rate.
Because these subscribers use conventional access rather than shared-media access through the DPU-C, the ASCII identifier (TLV 0x03) does not begin with the # character. In this case, the value is just the ASCII equivalent of the binary value conveyed in TLV 0x06.

Figure 13 on page 176 shows the five-level CoS hierarchy that corresponds to the CuTTB topology in Figure 12 on page 174.

Figure 13: CoS Hierarchy for CuTTB Topology



The following stanzas are part of the use case configuration that creates the Level 2 and Level 3 interface sets. The **stacked-interface-set** statement sets the Level 2 interface set to the `$junos-aggregation-interface-set-name` predefined variable. The stanza also specifies the Level 3 interface set as `$junos-interface-set-name`. It establishes the Level 2 set as the parent of the Level 3 set.

```
dynamic-profiles test-prof
  interfaces {
    stacked-interface-set {
      interface-set "$junos-aggregation-interface-set-name" {
        interface-set $junos-interface-set-name;
      }
    }
  }
}
```

The **predefined-variable-defaults** stanza uses variable expressions that set conditions to establish the names of the Level 2 and Level 3 interface sets. The default values are used only when RADIUS does not supply values for `$junos-aggregation-interface-set-name` and `$junos-interface-set-name`.

```
dynamic-profiles test-prof
```

```

predefined-variable-defaults {
    aggregation-interface-set-name equals "$junos-phy-ifd-underlying-intf-set-name";
    interface-set-name equals "ifZero($junos-default-interface-set-name, $junos-phy-ifd-interface-set-name)";
    default-interface-set-name equals "ifZero($junos-interface-set-name,
        ifNotZero($junos-aggregation-interface-set-name, $junos-aggregation-interface-set-name##'-default'))";
}
}

```

The following list describes the hierarchical CoS scheduler nodes for the CuTTB topology. It explains how the names of the interface sets are derived from predefined variables.

- Level 1 corresponds to the access-facing physical interface for all subscribers, xe-1/1/0.
- Level 2 corresponds to a parent interface set that has child interface sets as its members. The name of the interface set is supplied by the \$junos-aggregation-interface-set-name predefined variable in the dynamic profile.

- TEST-X2CTTB is the Level 2 interface set for all shared-media access subscribers. Its members are the Level 3 interface sets for residential subscribers 2 and 3 and for business subscriber 4.

TLV 0x03 includes the # character, which identifies the line as shared.

\$junos-aggregation-interface-set-name takes the value of TLV 0x03.

- xe-1/1/0-underlying is the Level 2 interface set for conventional access. Its member is the Level 3 interface set for residential subscribers 10 and 11.

TLV 0x03 does not include the # character and so does not identify a shared line.

\$junos-aggregation-interface-set-name is dynamically taken from

\$junos-phy-ifd-underlying-intf-set-name. The value of \$junos-phy-ifd-underlying-intf-set-name is simply the physical interface name with a suffix of “-underlying”.

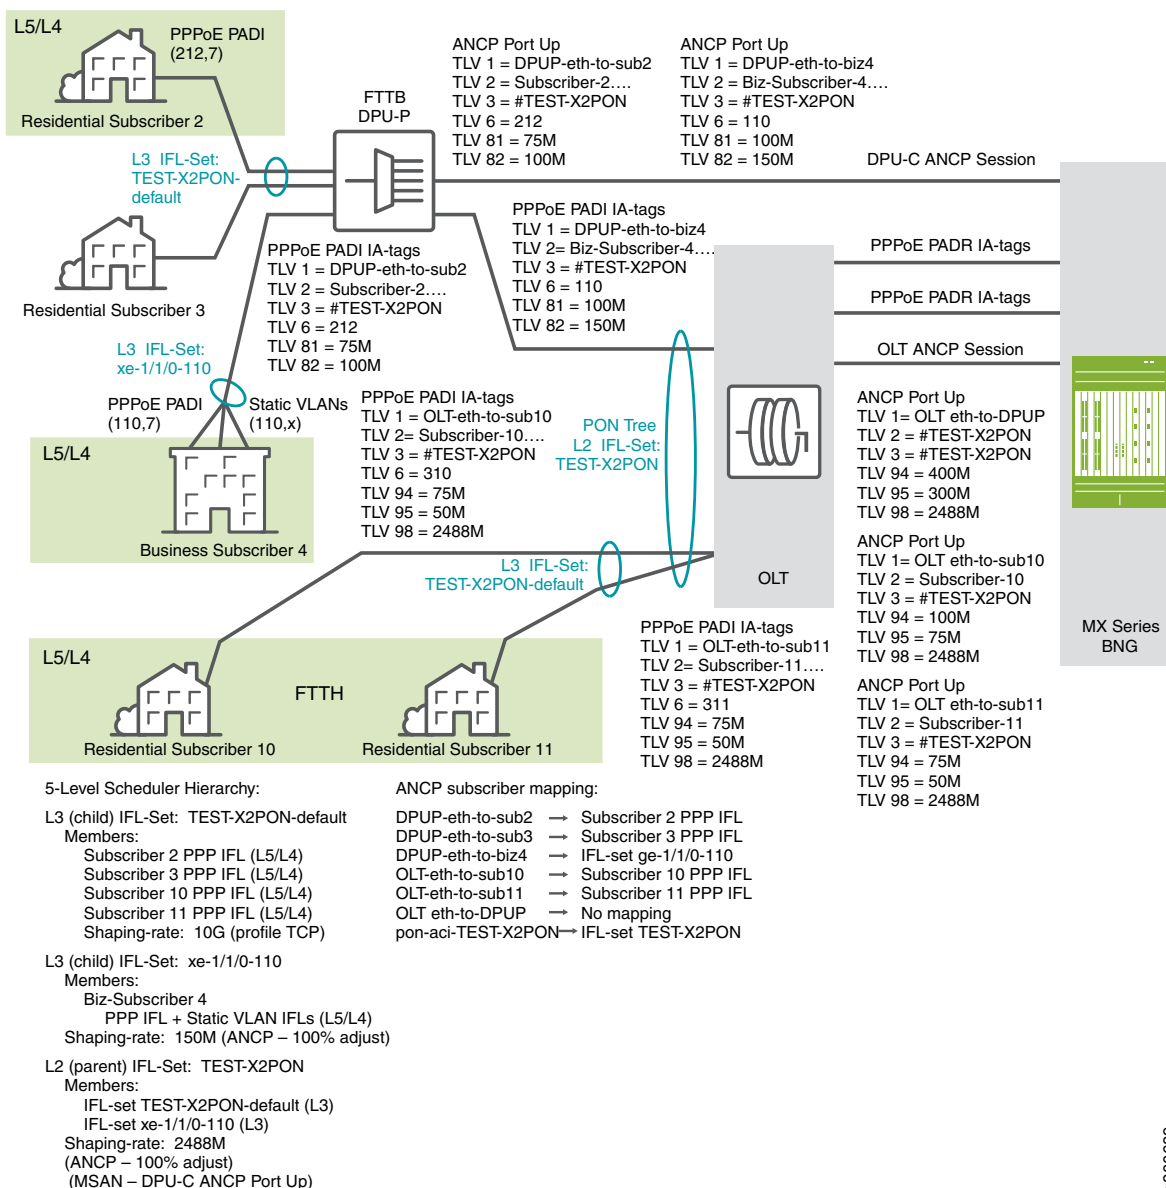
- Level 3 corresponds to a child interface set that has subscriber logical interfaces as its members. The name of the interface set is supplied by the \$junos-interface-set-name predefined variable in the dynamic profile.
 - TEST-X2CTTB-default is the Level 3 interface set for residential subscribers 2 and 3. These subscribers were identified as residential because the RADIUS server did not return VSA 26-4874-130, QoS-Set-Name. TLV 0x03 includes the # character, which identifies the line as shared. \$junos-interface-set-name is set to the value of \$junos-aggregation-interface-set-name with a suffix of “default”.
 - xe-1/1/0-110 is the Level 3 interface set for business subscriber 4. This subscriber was identified as business because the RADIUS server returned VSA 26-4874-130. TLV 0x03 includes the # character, which identifies the line as shared. \$junos-interface-set-name is set to the value of VSA 26-4874-130. The VSA value is a concatenation of the physical interface name (\$junos-phy-ifd-intf-set-name) and the outer VLAN tag.

- xe-1/1/0 is the Level 3 interface set for residential subscribers 10 and 11, which use conventional access. These subscribers were identified as residential because the RADIUS server did not return VSA 26-4874-130. \$junos-interface-set-name is set to the value of \$junos-phy-ifd-intf-set-name.
- Level 4 corresponds to the logical interface for individual subscribers. This includes PPPoE logical interfaces for residential and business subscribers, as well as static VLAN logical interfaces for business subscribers.
- Level 5 corresponds to the scheduling queue for each subscriber, regardless of subscriber type or access type. One or more queues are present per subscriber to provide subscriber services.

FTTB/FTTH Use Case Topology and CoS Hierarchy

[Figure 14 on page 179](#) shows a heterogeneous FTTB/FTTH topology that includes both shared -media (PON through a DPU-P) and conventional, (directly connected) access for PPPoE subscribers.

Figure 14: FTTB/FTTH CoS Hierarchy Example



This topology has the following subscribers:

- Two residential subscribers, 2 and 3, and a business subscriber, 4, have a shared-media access to the network through a DPU-P to the OLT and then to the BNG. These are FTTB subscribers.
- Two residential subscribers, 10 and 11, have conventional access to the network through the same OLT to the BNG. These are FTTH subscribers.

NOTE: All the FTTB and FTTH subscribers connect to the BNG by means of the same PON tree at the OLT.

- Residential subscriber 3 is not currently logged in.
- When residential subscriber 2 and business subscriber 4 log in:
 1. PPPoE sends a PADI message to the DPU-P that includes the outer VLAN tag for each.
 2. The DPU-P sends an ANCP Port Up message to the BNG for each subscriber. ANCP TLVs in the message identify the access line, the subscriber, the ASCII identifier for the access line, the VLAN outer tag as the binary identifier for the access line, the upstream rate, and the downstream rate.

The ASCII identifier (TLV 0x03) begins with the # character, signifying that the remainder of the value identifies the backhaul (PON tree) line. TLV 0x03 is the same for both subscribers, because they connect through the same PON tree.
 3. The DPU-P sends a PADI message for each subscriber to the OLT. The PADI conveys the PPPoE-IA tags that identify the same attributes as the ANCP TLVs.
 4. The OLT sends a PADR message with the PPPoE tags to the BNG. The OLT also opens an ANCP session to the BNG by sending an ANCP Port Up message for the OLT-to-DPU-P connection. The rates in TLVs 0x81 and 0x82 are the values for the OLT-to-DPU-P line, represented by the L2 interface set. In other words, these are the rates for the PON tree itself rather than the subscriber access lines. Although this use case example shows that TLV 0x03 value is also reported in TLV 0x02 to indicate the PON tree line, this is not a requirement for PON networks.

NOTE: The FTTB portion of this network connects G.fast DSL subscribers to the PON tree shared media backhaul. Consequently the DPU-P reports DSL TLVs for these subscribers rather than PON TLVs.

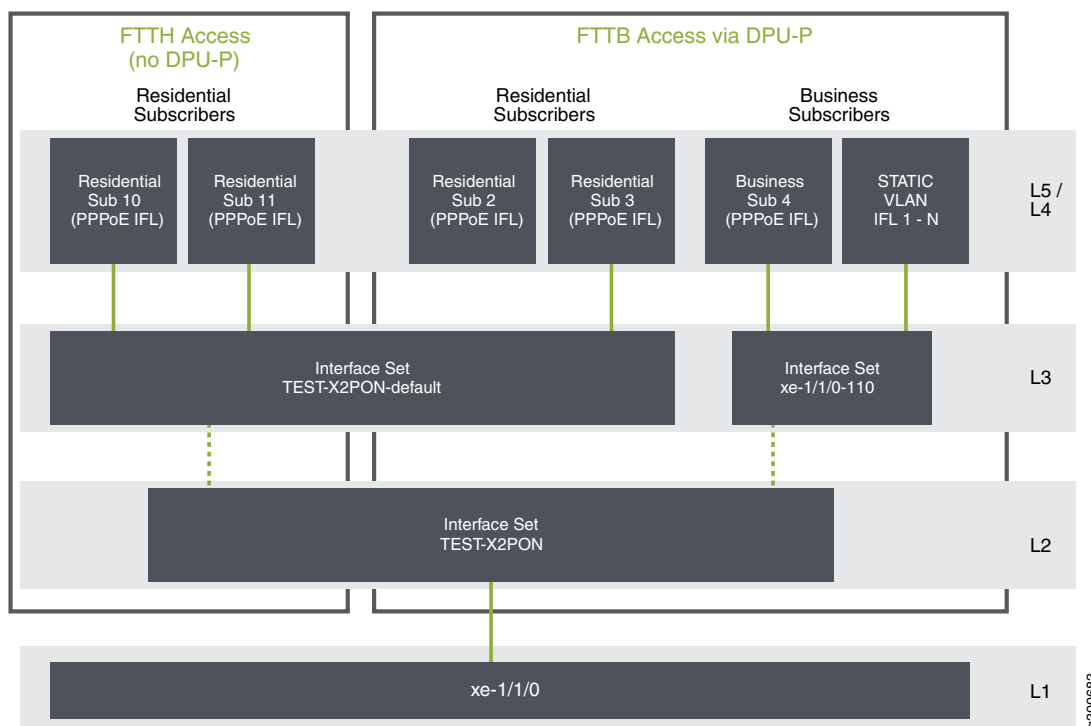
- When residential subscribers 10 and 11 log in:
 1. PPPoE sends a PADI message for each subscriber to the OLT. The PADI conveys the PPPoE-IA tags for the individual subscriber access lines.
 2. The OLT sends a PADR message with the PPPoE tags for each subscriber line to the BNG.
 3. The OLT also sends an ANCP Port Up message to the BNG for each subscriber. ANCP TLVs in the message identify the access line, the subscriber, the ASCII identifier for the access line, the VLAN outer tag as the binary identifier for the access line, the upstream rate, and the downstream rate.

The ASCII identifier (TLV 0x03) begins with the # character, signifying that the remainder of the value identifies the backhaul (fiber PON tree, shared media) line. TLV 0x03 is the same for both subscribers, because they connect through the same DPU-P.

Because subscribers 10 and 11 connect to the same PON tree as the FTTB subscribers, the ASCII identifier (TLV 0x03) also begins with the # character, signifying that the remainder of the value identifies the backhaul (fiber PON tree, shared media) line. TLV 0x03 is the same for both subscribers.

Figure 15 on page 181 shows the five-level CoS hierarchy that corresponds to the FTTB/FTTH topology in Figure 14 on page 179.

Figure 15: CoS Hierarchy for FTTB/FTTH Topology



The following stanzas are part of the use case configuration that creates the Level 2 and Level 3 interface sets. The **stacked-interface-set** statement sets the Level 2 interface set to the `$junos-aggregation-interface-set-name` predefined variable. The stanza also specifies the Level 3 interface set as `$junos-interface-set-name`. It establishes the Level 2 set as the parent of the Level 3 set.

```
dynamic-profiles test-prof
  interfaces {
    stacked-interface-set {
      interface-set "$junos-aggregation-interface-set-name" {
        interface-set $junos-interface-set-name;
      }
    }
  }
```

```
}
```

The **predefined-variable-defaults** stanza uses variable expressions that set conditions to establish the names of the Level 2 and Level 3 interface sets. The default values are used only when RADIUS does not supply values for `$junos-aggregation-interface-set-name` and `$junos-interface-set-name`.

```
dynamic-profiles test-prof
  predefined-variable-defaults {
    aggregation-interface-set-name equals "$junos-phy-ifd-underlying-intf-set-name";
    interface-set-name equals "ifZero($junos-default-interface-set-name, $junos-phy-ifd-interface-set-name)";
    default-interface-set-name equals "ifZero($junos-interface-set-name,
      ifNotZero($junos-aggregation-interface-set-name, $junos-aggregation-interface-set-name##'-default'))";
  }
}
```

The following list describes the hierarchical CoS scheduler nodes for the FTTB/FTTH topology. It explains how the names of the interface sets are derived from predefined variables.

- Level 1 corresponds to the access-facing physical interface for all subscribers, xe-1/1/0.
- Level 2 corresponds to a parent interface set that has child interface sets as its members. The name of the interface set is supplied by the `$junos-aggregation-interface-set-name` predefined variable in the dynamic profile.

TEST-X2PON is the Level 2 interface set for all PON subscribers, both conventional access and DPU-P access. Its members are the Level 3 interface set for the FTTB/FTTH residential subscribers and the Level 3 interface set for business subscriber 4. TLV 0x03 includes the # character, which identifies the PON tree line as the backhaul. `$junos-aggregation-interface-set-name` takes the value of TLV 0x03.

- Level 3 corresponds to an interface set that has subscriber logical interfaces as its members.
 - TEST-X2PON-default is the Level 3 interface set for FTTB residential subscribers 2 and 3, as well as FTTH residential subscribers 10 and 11. These subscribers all use the same PON tree and therefore are included in the same interface set.

These subscribers were identified as residential because the RADIUS server did not return VSA 26-4874-130, QoS-Set-Name. TLV 0x03 includes the # character, which identifies the PON tree line as the backhaul. `$junos-interface-set-name` is set to the value of `$junos-aggregation-interface-set-name` with a suffix of "default".

- xe-1/1/0-110 is the Level 3 interface set for business subscriber 4, which uses shared-media access.

This subscriber was identified as business because the RADIUS server returned VSA 26-4874-130.

TLV 0x03 includes the # character, which identifies the PON tree line as the backhaul.

`$junos-interface-set-name` is set to the value of VSA 26-4874-130. The VSA value is a concatenation of the physical interface name (`$junos-phy-ifd-intf-set-name`) and the outer VLAN tag.

- Level 4 corresponds to the logical interface for individual subscribers. This includes PPPoE logical interfaces for residential and business subscribers, as well as static VLAN logical interfaces for business subscribers.
- Level 5 corresponds to the scheduling queue for each subscriber, regardless of subscriber type or access type. One or more queues are present per subscriber to provide subscriber services.

Automatic Creation of Business Subscriber Interface Sets

For business subscribers in an access network, four-level scheduler hierarchies use static interface sets to represent the subscriber access line. The members of the interface set are static VLAN logical interfaces. This configuration is performed by Extensible Subscriber Services Manager (ESSM) operation scripts (op-scripts).

The op-scripts base the name on the outer VLAN tag of the subscriber interface, because the tag is unique per subscriber. The interface set name is in the format *physical_interface_name-outer_vlan_tag*. For example, an Ethernet interface ge-1/1/0, with a dual-tagged VLAN interface that has an outer tag of 111, results in an interface set name of ge-1/1/0-111. This format is the same as that used by the `$junos-svlan-interface-set-name` predefined variable.

In five-level scheduler hierarchies for business subscribers, each business session includes a dynamic PPPoE control session (and thus a dynamic PPPoE logical interface) and two or more static business VLAN logical interfaces. These interfaces need to be shaped as an aggregate in an interface set. The dynamic logical interfaces cannot be assigned to a static interface set. This means that this deployment design requires dynamic interface sets for logical intermediate (Level 3) CoS nodes to accommodate both the dynamic PPPoE logical interfaces and the static interfaces.

BEST PRACTICE: We recommend that you use dynamic interface sets to provide a uniform solution for both four-level and five-level hierarchies. This method ensures that all the logical interfaces, both dynamic and static, are members of the same interface set. This is not a requirement. You can continue to configure only static interface sets for business subscribers in four-level hierarchies.

The op-scripts need to reference the business subscriber dynamic interface set name during subscriber configuration. This means that the format of the dynamic interface set name must be the same format that the script uses for static interface sets. The interface set name is provided by the RADIUS server during subscriber authentication, because the server has to determine whether the subscriber logging in is a business subscriber or a residential subscriber. This means that you have to configure your RADIUS software to specify the interface set for each subscriber. This requirement adds initial and maintenance configuration overhead to your operations, especially as your networks scales to higher numbers of subscribers.

Starting in Junos OS Release 19.3R1, you can configure the BNG to dynamically create the interface set name and propose that name to the RADIUS server in the Access-Request message for the subscriber. This method reduces the complexity of the RADIUS configuration, because you avoid having to configure your RADIUS software to specify interface sets for each subscriber. To enable dynamic creation of the interface set name for business subscribers, use the **source-interface-set-at-login svlan** statement at the **[edit protocols ppp-service]** hierarchy level.

The interface set name that the BNG proposes is carried by the Juniper Networks VSA, Qos-Set-Name (26-130) in the RADIUS Access-Request message. The set name consists of the name of access-facing physical interface appended with the VLAN tag. This is the same format that is used by the op scripts:

- The outer vlan tag is used for a dual-tagged VLAN. For a business subscriber on xe-1/1/0 with VLAN tags (110,7), the name has this format:
xe-1/1/0-110
- The lone vlan tag is used for a single-tagged VLAN. The single-tagged VLAN is used when the CPE device connects directly to the access node. For a business subscriber on xe-2/2/1 with VLAN tag (33), the name has this format:
xe-2/2/1-33

When the subscriber logs in, the RADIUS server evaluates the Access-Request and determines whether the subscriber is business or residential:

- When the RADIUS server determines that the subscriber is a business subscriber, it returns the VSA with the name in the Access-Accept message to the BNG, where the name is used to create a dynamic interface set for the business subscriber.
- If the RADIUS server determines during authentication that the subscriber is residential, then the server does not return the VSA in the Access-Accept message. In this case, the dynamic PPPoE IFL is added to a default dynamic interface set to conserve L3 CoS nodes for a five-level hierarchy or L2 CoS nodes for a four-level hierarchy. The dynamic interface set for residential subscribers always resolves to the default interface set. The default dynamic interface set is determined by how you configure the **predefined-variable-defaults** statement with expressions in the dynamic profile. See [“Dynamic Level 2 and Level 3 Interface Set Naming with Predefined Variables” on page 185](#) for information about configured the defaults.

How to Configure the Automatic Creation of Business Subscriber Interface Sets

In heterogeneous access networks, you can reduce some of the complexity of your RADIUS configuration by having PPP on the BNG dynamically create the interface set name for business subscribers and propose that name to the RADIUS server in the Access-Request message for the subscriber. This method reduces complexity because you do not have to configure all the possible interface set names on the RADIUS server. The proposed name is carried by the Qos-Set-Name VSA (26-4874–130).

If the server determines that the subscriber is a business subscriber, it returns the name in the Access-Accept message to the BNG. PPP on the BNG then uses the name to create a dynamic interface set for the business subscriber. This interface set is for an intermediate CoS node; for example, Level 3 in a five-level hierarchy. This interface set includes the business subscriber PPPoE IFL and the static VLAN IFLs created by ESSMD op-scripts. It is the child interface set of the Level 2 parent interface set.

For information about how the interface set names are formed, see [“Automatic Creation of Business Subscriber Interface Sets” on page 183](#).

If the RADIUS server determines that the subscriber is residential, then the server does not return the VSA in the Access-Accept message. In this case, the dynamic PPPoE IFL is added to a default dynamic interface set.

To configure dynamic creation of business subscriber interface sets with the same format as `$junos-svlan-interface-set-name`:

- Enable PPP to dynamic creation.

```
[edit protocols ppp-service]
user@host# set source-interface-set-at-loginsvlan
```

Dynamic Level 2 and Level 3 Interface Set Naming with Predefined Variables

In heterogeneous access networks, Juniper Networks predefined variables supply the names of the interface sets for the Level 2 and Level 3 CoS nodes:

- Level 2—`$junos-aggregation-interface-set-name`
- Level 3—`$junos-interface-set-name`

You specify these variables in the dynamic client profile at the [\[edit dynamic-profiles *profile-name* interfaces\]](#) hierarchy level as follows:

```
stacked-interface-set {
  interface-set "$junos-aggregation-interface-set-name" {
    interface-set "$junos-interface-set-name";
  }
}
```

These interfaces are said to be stacked. Level 2 is the parent interface set and Level 3 is the child interface set.

You can optionally configure default values for the predefined variables. The default value must be appropriate to the variable, such as an integer or an alphanumeric string. The Junos OS uses the default value when the variable is not resolved, meaning that it does not have a value. The predefined variable

might not be resolved for several reasons, depending on the access type (conventional or shared-media) and the subscriber type (residential or business), such as the following:

- The Access-Aggregation-Circuit-Id-ASCII TLV (0x03) is not present or does not include the # character that indicates it carries the backhaul identifier).
- The external RADIUS server does not return the QoS-Set-Name VSA (26-4874-130).

Starting in Junos OS Release 19.3R1, you can configure the default value of a predefined variable to be another predefined variable by using a variable expression. In earlier releases, the default value must be fixed; it cannot be a variable.

NOTE: Expressions are typically configured for user-defined variables and dynamic service profiles. See *Using Variable Expressions in User-Defined Variables* for more information.

When you use a variable expression, you are setting up a condition that determines the default value of the predefined variable. The value of the default is different when the condition is matched than when it is not matched. This capability enables you to configure a single dynamic client profile for a heterogeneous network. The profile can instantiate the proper interface sets for business subscribers and residential subscribers on both conventional access lines and shared-media access lines.

In dynamic client profiles, you can configure variable expressions that use any of the following:

- **equals**—Assigns a predefined variable or expression as the default value.
- **ifNotZero(*parameter-1*, *parameter-2*)**—Sets a condition to be matched. Assigns the value from *parameter-2* as the default value only when *parameter-1* is nonzero, meaning that the parameter resolved to some value.
- **ifZero(*parameter-1*, *parameter-2*)**—Sets a condition to be matched. Assigns the value from *parameter-2* as the default value only when *parameter-1* is zero, meaning that the parameter did not resolve to any value. If *parameter-1* did resolve to a value (therefore it is not zero), then the value from *parameter-1* is assigned as the default.

You can also nest expressions, which provides additional conditions for setting the variable value. For a heterogeneous network, you use the following expressions to determine the name for the Level 2 and Level 3 CoS nodes:

```
dynamic-profiles name {
  predefined-variable-defaults {
    interface-set-name equals "ifZero($junos-default-interface-set-name, $junos-phy-ifd-interface-set-name)";
    default-interface-set-name equals "ifZero($junos-interface-set-name,
      ifNotZero($junos-aggregation-interface-set-name, $junos-aggregation-interface-set-name##'-default'))";
  }
}
```

```
}

```

The following sections explain how to evaluate each of these expressions.

Predefined Variable Default for the Level 2 Node Interface Set

The following definition simply assigns a predefined variable as the default value for \$junos-aggregation-interface-set-name:

```
aggregation-interface-set-name equals "$junos-phy-ifd-underlying-intf-set-name"

```

The expression has no conditions to evaluate. The \$junos-phy-ifd-underlying-intf-set-name predefined variable has the format *physical-interface-name*-underlying. For example, if the physical interface is xe-1/1/0, then \$junos-phy-ifd-underlying-intf-set-name resolves to xe-1/1/0-underlying. That becomes the default value for \$junos-aggregation-interface-set-name:

\$junos-aggregation-interface-set-name = \$junos-phy-ifd-underlying-intf-set-name = xe-1/1/0-underlying

The default value is not used when \$junos-aggregation-interface-set-name is already resolved. If the Access-Aggregation-Circuit-ID-ASCII attribute (TLV 0x03) begins with a # character (the backhaul identifier), then the variable takes the value of the remainder of the string after the # character. It is therefore resolved and the default is not used.

The following table shows the value of \$junos-aggregation-interface-set-name when TLV 0x03 identifies the backhaul node and when it is not present. The physical interface is xe-1/1/0.

TLV 0x03 (Access Type)	\$junos-aggregation-interface-set-name
#TEST-X2PON (DPU-C/DPU-P)	TEST-X2PON
Not present in PPPoE-IA tags (Conventional)	xe-1/1/0-underlying

Predefined Variable Default for the Level 3 Node Interface Set

You have to use multiple expressions to provide a default value for \$junos-interface-set-name:

```
interface-set-name equals "ifZero($junos-default-interface-set-name, $junos-phy-ifd-interface-set-name)";
default-interface-set-name equals "ifZero($junos-interface-set-name,
    ifNotZero($junos-aggregation-interface-set-name, $junos-aggregation-interface-set-name##'-default'))";

```

1. The first expression means that it has to check whether \$junos-default-interface-set-name is resolved.

```
interface-set-name equals "ifZero($junos-default-interface-set-name, $junos-phy-ifd-interface-set-name)";
```

- If it is not resolved, then the default value for \$junos-interface-set-name is set to the value of \$junos-phy-ifd-interface-set-name:

```
$junos-interface-set-name = $junos-phy-ifd-interface-set-name
```

- If it is resolved, then the default value for \$junos-interface-set-name is set to the resolved value of \$junos-default-interface-set-name:

```
$junos-interface-set-name = $junos-default-interface-set-name
```

2. The value of \$junos-default-interface-set-name is determined by a nested expression.

```
default-interface-set-name equals "ifZero($junos-interface-set-name,  
ifNotZero($junos-aggregation-interface-set-name, $junos-aggregation-interface-set-name##'-default'))";
```

- a. If \$junos-interface-set-name is not resolved, then \$junos-interface-set-name is set to the result of the nested expression (ifNotZero). However, the predefined variable defaults are used only if \$junos-interface-set-name is not resolved. Consequently, the expression must reduce to this:

```
default-interface-set-name equals  
"ifNotZero($junos-aggregation-interface-set-name,  
$junos-aggregation-interface-set-name##'-default')"
```

- b. The ifNotZero expression is solved by evaluating whether \$junos-aggregation-interface-set-name is resolved. \$junos-aggregation-interface-set-name is resolved only when TLV 0x03 includes the backhaul identifier (#).

- If \$junos-aggregation-interface-set-name is resolved, then **-default** is appended to that name and that becomes the default value for \$junos-default-interface-set-name:

```
$junos-default-interface-set-name = $junos-aggregation-interface-set-name+ "-default"
```

- If \$junos-aggregation-interface-set-name is not resolved, then \$junos-default-interface-set-name is also not resolved.

3. Now the value for \$junos-interface-set-name can be determined:

```
interface-set-name equals "ifZero($junos-default-interface-set-name, $junos-phy-ifd-interface-set-name)";
```

- If \$junos-default-interface-set-name is resolved, then that is also the value of \$junos-interface-set:

```
$junos-interface-set-name = $junos-default-interface-set-name =  
$junos-aggregation-interface-set-name+ "-default"
```

- If \$junos-default-interface-set-name is not resolved, then:

`$junos-interface-set-name = $junos-phy-ifd-interface-set-name`

The following table shows the possible values of the predefined variables based on the expressions described above. It can be helpful to refer to the figures and text in [“CoS Node Shaping in Four-Level and Five-Level Heterogeneous Networks” on page 168](#), [“CuTTB Use Case Topology and CoS Hierarchy” on page 173](#), and [“FTTB/FTTH Use Case Topology and CoS Hierarchy” on page 178](#)

TLV 0x03 (Access Type)	VSA 26-4874-130 (Subscriber Type)	<code>\$junos-phy-ifd-interface-set-name</code>	<code>\$junos-default-interface-set-name</code>	<code>\$junos-int</code>
#TEST-X2PON (DPU-C/DPU-P)	Not returned (Residential)	xe-1/1/0	Not resolved	TEST-X2PO
#TEST-X2PON (DPU-C/DPU-P)	Returned as xe-1/1/0 (Business)	xe-1/1/0	xe-1/1/0	xe-1/1/0-1
Not present in PPPoE-IA tags (Conventional)	Not returned (Residential)	xe-1/1/0	Not resolved	xe-1/1/0
Not present in PPPoE-IA tags (Conventional)	Returned as xe-1/1/0 (Business)	xe-1/1/0	xe-1/1/0	xe-1/1/0-1

Release History Table

Release	Description
19.3R1	Starting in Junos OS Release 19.3R1, you can configure the BNG to dynamically create the interface set name and propose that name to the RADIUS server in the Access-Request message for the subscriber.
19.3R1	Starting in Junos OS Release 19.3R1, you can configure the default value of a predefined variable to be another predefined variable by using a variable expression.

RELATED DOCUMENTATION

ANCP Agent and AAA

ANCP Agent Traffic Shaping and CoS

DSL Forum Vendor-Specific Attributes

[OLT Migration to Using PON TLVs Instead of DSL TLVs | 190](#)

CoS for Subscriber Access Overview

Hierarchical Class of Service for Subscriber Management Overview

Predefined Variables in Dynamic Profiles

OLT Migration to Using PON TLVs Instead of DSL TLVs

IN THIS SECTION

- [Support for OLT Migration to PON TLVs | 190](#)
- [How to Configure Preference for DSL or PON TLVs When an OLT Sends Both | 191](#)

Support for OLT Migration to PON TLVs

Before the introduction of PON-specific TLVs (see the Internet draft, *Access Extensions for the Access Node Control Protocol*), OLTs passed information about PON access line rates in the ANCP DSL TLVs. The DSL TLVs are *overloaded* with the PON data. In this situation, the access line type is classified as OTHER in the DSL-Type TLV (0x91). The raw PON line rates are conveyed in the Actual-Net-Data-Rate-Upstream TLV (0x81) and the Actual-Net-Data-Rate-Downstream TLV (0x82). For example, this is the behavior used for FTTH networks.

With the availability of the PON TLVs, OLTs provide PON access line information in one of the following ways:

- Only in DSL TLVs—OLTs that do not support the PON TLVs continue to send PON rates in DSL TLVs. This is the same situation as before the introduction of PON TLVs.
- Only in PON TLVs—OLTs that support the PON TLVs might use them exclusively to convey attributes for the subscriber access line and PON tree.
- In both DSL and PON TLVs—OLTs that support the PON TLVs might send them and also, redundantly, overload the DSL TLVs with the PON rates. This behavior enables the OLT to successfully provide rate information both to BNGs that support PON TLVs and BNGs that do not. The OLT is expected to provide both types of TLVs in both the ANCP port status messages and the PPPoE-IA tags.

The following expectations apply to OLTs that provide both DSL TLVs and PON TLVs:

- The OLT uses both types of TLVs for both ANCP port status messages and PPPoE-IA tags.
- The ANCP line attribute TLVs carry only the TLVs that match that line:
 - The DSL-Line-Attributes TLV (0x04) carries only DSL TLVs, including G.fast TLVs.
 - The PON-Access-Line-Attributes TLV (0x12) carries only PON TLVs.
 - The Access-Loop Encapsulation TLV (0x90) is independent of the transport method. The ANCP agent accepts and saves 0x90 when it is present in either the PON-Access-Line-Attributes TLV (0x12) or in PPPoE-IA tags.
- The access-line identification attributes are common to both DSL and PON access lines and convey the same information. This set of attributes consists of the following:
 - Access-Loop-Circuit-ID (0x01)
 - Access-Loop-Remote-ID (0x02)
 - Access-Aggregation-Circuit-ID-ASCII(0x03)
 - Access-Aggregation-Circuit-ID-Binary (0x04)

When the OLT sends both types of TLVs, the router accepts, saves, and process only one type or the other, according to a preference setting. The router saves or discards the corresponding line attribute that carries the TLVs for that line type. By default, the router prefers the PON TLVs over the DSL TLVs. This means that the router accepts the PON-Line-Attributes TLV (0x12) and discards the DSL-Line-Attributes-TLV (0x04).

You can change the preference with the **preference (dsl | pon)** option at the **[edit system access-line attributes]** hierarchy level. For example, when the PON TLVs are unreliable, perhaps because of an OLT issue, you might configure the router prefer the DSL TLVs for improved reliability.

How to Configure Preference for DSL or PON TLVs When an OLT Sends Both

You can configure which set of TLVs is saved and processed when the OLT sends both DSL TLVs and PON TLVs in ANCP port status messages or in PPPoE-IA tags. The TLVs that you do not select are discarded. OLTs provide PON access line information in one of the following ways:

- Only in DSL TLVs—OLTs that do not support the PON TLVs continue to use the method available before PON TLVs were available. These OLTs send PON rates in DSL TLVs, overloading the DSL TLVs with the PON information. In this case, the DSL-Type TLV (0x91) is set to OTHER and PON rates for the subscriber access line are presented in the Actual-Net-Data-Rate-Upstream TLV (0x81) and the Actual-Net-Data-Rate-Downstream TLV (0x82).
- Only in PON TLVs—OLTs that support the PON TLVs might use them exclusively to convey attributes for the subscriber access line and PON tree.

- In both DSL and PON TLVs—OLTs that support the PON TLVs might send them and also, redundantly, overload the DSL TLVs with the PON rates. This behavior enables the OLT to successfully provide rate information both to BNGs that support PON TLVs and BNGs that do not. The OLT is expected to provide both types of TLVs in both the ANCP port status messages and the PPPoE-IA tags.

NOTE: The default preference on the router is to accept and process PON TLVs when both are received. This means that you typically use the **preference** option to prioritize DSL TLVs over PON TLVs. For example, if the PON TLVs are unreliable, perhaps because of an OLT issue, you can prioritize the DSL TLVs for reliability.

To select the type of access line TLVs to accept:

- Configure the preference.

```
[edit system access-line]
user@host# set preference (dsl | pon)
```

RELATED DOCUMENTATION

DSL Forum Vendor-Specific Attributes

DSL Forum VSAs Support in AAA Access and Accounting Messages for Junos OS

[Five-Level and Four-Level Heterogeneous Networks](#) | 168

5

PART

Configuration Statements and Operational Commands

Configuration Statements | **194**

Operational Commands | **508**

Configuration Statements

IN THIS CHAPTER

- `accept` | 198
- `accept-out-of-band` | 200
- `access-profile` | 201
- `access-profile (Dynamic VLAN)` | 203
- `access-profile (Dynamic Stacked VLAN)` | 204
- `active-server-group` | 205
- `address` | 207
- `address-assignment (Address-Assignment Pools)` | 211
- `adjacency-loss-hold-time (ANCP)` | 214
- `ancp` | 215
- `authentication` | 217
- `authentication (DHCP Local Server)` | 219
- `authentication (DHCP Relay Agent)` | 221
- `authentication-order` | 223
- `auto-configure` | 225
- `auto-configure-trigger interface (ANCP)` | 227
- `backup-on-failure (Accounting Options)` | 228
- `circuit-id (VLAN Authentication Username)` | 229
- `cleanup-interval (Accounting Options)` | 230
- `compress (Accounting Options)` | 231
- `connectivity-type` | 232
- `core-facing` | 233
- `demux0 (Dynamic Interface)` | 234
- `demux-options (Dynamic Interface)` | 236
- `demux-source (Dynamic IP Demux Interface)` | 237
- `demux-source (Dynamic Underlying Interface)` | 238
- `demux-source (Underlying Interface)` | 239
- `dhcp-attributes (Address-Assignment Pools)` | 240

- [dhcp-local-server | 247](#)
- [dhcp-relay | 264](#)
- [dhcpv6 \(DHCP Local Server\) | 283](#)
- [dynamic-profile \(DHCP Local Server\) | 292](#)
- [dynamic-profile \(DHCP Relay Agent\) | 294](#)
- [dynamic-profile \(Dynamic PPPoE\) | 296](#)
- [dynamic-profile \(Stacked VLAN\) | 298](#)
- [dynamic-profile \(VLAN\) | 299](#)
- [dynamic-profiles | 300](#)
- [egress-stats \(Flat-File Accounting Options\) | 314](#)
- [encapsulation \(Dynamic Interfaces\) | 316](#)
- [exclude \(RADIUS Attributes\) | 319](#)
- [family \(Address-Assignment Pools\) | 327](#)
- [family \(Dynamic Demux Interface\) | 329](#)
- [family \(Dynamic PPPoE\) | 331](#)
- [family \(Dynamic Standard Interface\) | 333](#)
- [fields \(Flat-File Accounting Options\) | 336](#)
- [file \(Flat-File Accounting Options\) | 339](#)
- [flat-file-profile \(Accounting Options\) | 340](#)
- [flat-file-profile \(Extensible Subscriber Services\) | 342](#)
- [flexible-vlan-tagging | 343](#)
- [format \(Flat-File Accounting Options\) | 344](#)
- [forwarding-options | 345](#)
- [general-param \(Flat-File Accounting Options\) | 348](#)
- [group \(DHCP Local Server\) | 350](#)
- [group \(DHCP Relay Agent\) | 356](#)
- [ingress-stats \(Flat-File Accounting Options\) | 362](#)
- [inner-vlan-id \(Dynamic VLANs\) | 364](#)
- [inner-vlan-id-swap-ranges | 365](#)
- [input-vlan-map \(Dynamic Interfaces\) | 366](#)
- [instance-role | 367](#)
- [instance-type | 369](#)
- [interface \(DHCP Local Server\) | 372](#)
- [interface \(DHCP Relay Agent\) | 375](#)

- interface (Dynamic Routing Instances) | 378
- interface (Routing Instances) | 379
- interface-mac-limit (VPLS) | 380
- interfaces (Static and Dynamic Subscribers) | 382
- interval (Flat-File Accounting Options) | 389
- ip-address-first | 390
- keepalives (Dynamic Profiles) | 391
- l2-stats (Flat-File Accounting Options) | 393
- mac-validate (Dynamic IP Demux Interface) | 394
- multicast-replication | 395
- neighbor (Define ANCP) | 397
- no-local-switching | 398
- no-tunnel-services | 399
- overall-packet (Flat-File Accounting Options) | 401
- output-vlan-map (Dynamic Interfaces) | 403
- pap (Dynamic PPP) | 404
- pool (Address-Assignment Pools) | 405
- pool-match-order | 407
- pop (Dynamic VLANs) | 408
- pppoe-options (Dynamic PPPoE) | 409
- pppoe-underlying-options (Static and Dynamic Subscribers) | 410
- ppp-options (Dynamic PPP) | 411
- prefix (Address-Assignment Pools) | 413
- profile (Access) | 414
- protocols | 421
- proxy-arp | 425
- proxy-arp (Dynamic Profiles) | 426
- push (Dynamic VLANs) | 427
- push-backup-to-master (Accounting Options) | 428
- radius (Access Profile) | 429
- radius-server | 433
- range (Address-Assignment Pools) | 439
- ranges (Dynamic VLAN) | 440
- remote-id (VLAN Authentication Username) | 441

- route-distinguisher | 442
- routing-instances (Dynamic Profiles) | 446
- routing-instances (Multiple Routing Entities) | 448
- schema-version (Flat-File Accounting Options) | 450
- secret (RADIUS) | 451
- server (Dynamic PPPoE) | 452
- server-group | 453
- site (VPLS Multihoming for FEC 128) | 455
- site-identifier (VPLS) | 456
- site-range | 457
- stacked-vlan-ranges | 458
- stacked-vlan-tagging | 459
- traceoptions (DHCP) | 460
- underlying-interface (demux0) | 463
- underlying-interface (Dynamic PPPoE) | 465
- unit | 466
- unit (Dynamic Demux Interface) | 477
- unit (Dynamic Profiles Standard Interface) | 479
- unnumbered-address (Dynamic PPPoE) | 483
- unnumbered-address (Dynamic Profiles) | 484
- unnumbered-address (Ethernet) | 486
- username-include (Interfaces) | 487
- user-prefix (DHCP Local Server) | 489
- vlan-id (Dynamic VLANs) | 491
- vlan-model | 492
- vlan-ranges | 493
- vlan-tags | 495
- vlan-tags (Stacked VLAN Tags) | 497
- vpls (Routing Instance) | 499
- vrf-export | 502
- vrf-import | 504
- vrf-target | 506

accept

Syntax

```
accept (any | dhcp-v4 | dhcp-v6 | inet | inet6 | pppoe);
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure stacked-vlan-ranges dynamic-profile profile-name],  
[edit interfaces interface-name auto-configure vlan-ranges dynamic-profile profile-name]
```

Release Information

Statement introduced in Junos OS Release 9.5.

dhcp-v4 option added in Junos OS Release 10.0.

dhcp-v6, **inet6** and **pppoe** options added in Junos OS Release 10.2.

any option added in Junos OS Release 10.4.

Description

Specify the type of VLAN Ethernet packet accepted by an interface that is associated with a VLAN dynamic profile or stacked VLAN dynamic profile.

Options

any—Any packet type. Specifies that any incoming packets trigger the dynamic creation of a VLAN with properties determined by the auto-configure interface configuration stanza and associated profile attributes. This option is used when configuring wholesaling in a Layer 2 network.

dhcp-v4—IPv4 DHCP packet type. Specifies that incoming IPv4 DHCP discover packets trigger the dynamic creation of a VLAN with properties determined by the auto-configure interface configuration stanza and associated profile attributes

NOTE: The DHCP-specific **mac-address** and **option-82** options are rejected if the **accept** statement is not set to **dhcp-v4**.

dhcp-v6—IPv6 DHCP packet type. Specifies that incoming IPv6 DHCP discover packets trigger the dynamic creation of a VLAN with properties determined by the auto-configure interface configuration stanza and associated profile attributes.

inet—IPv4 Ethernet and ARP packet type.

inet6—IPv6 Ethernet packet type.

pppoe—Point-to-Point Protocol over Ethernet packet type.

NOTE: The **pppoe** VLAN Ethernet packet type option is supported only for MPC/MIC interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs

Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs

[Configuring VLAN Interfaces for the Layer 2 Wholesale Solution | 97](#)

Configuring Subscriber Packet Types to Trigger VLAN Authentication

accept-out-of-band

Syntax

```
accept-out-of-band protocol;
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges dynamic-profile profile-name]
```

Release Information

Statement introduced in Junos OS Release 16.1R4.

Description

Configure the protocol for which packets are accepted as out-of-band traffic to trigger instantiation or deletion of autosensed dynamic VLANs.

NOTE: A given physical interface can support VLANs created by either conventional packet-triggering or out-of-band triggering, but not both at the same time.

Options

protocol—Out-of-band protocol. The following out-of-band protocol is supported:

- **ancp**—ANCP Port Up and Port Down messages trigger instantiation and deletion, respectively, of autosensed, dynamic VLAN.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Out-of-Band ANCP Messages to Trigger Dynamic VLAN Instantiation | 141](#)

[Layer 2 Wholesale with ANCP-Triggered VLANs Overview | 116](#)

access-profile

Syntax

```
access-profile profile-name;
```

Hierarchy Level

```
[edit],
[edit forwarding-options dhcp-relay]
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name],
[edit forwarding-options dhcp-relay group group-name]
[edit forwarding-options dhcp-relay dhcpv6]
[edit forwarding-options dhcp-relay dhcpv6 group group-name]
[edit logical-systems logical-system-name routing-instances routing-instance-name]
[edit interfaces interface-name auto-configure vlan-ranges],
[edit interfaces interface-name auto-configure stacked-vlan-ranges],
[edit routing-instances routing-instances-name]
[edit system services dhcp-local-server]
[edit system services dhcp-local-server group group-name]
[edit system services dhcp-local-server dhcpv6]
[edit system services dhcp-local-server dhcpv6 group group-name]
[edit system services dhcp-local-server dual-stack-group dual-stack-group-name]
```

Release Information

Statement introduced in Junos OS Release 9.1.

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

Description

After you have created the access profile that specifies authentication and accounting parameters, you must specify where the profile is used. Authentication and accounting will not run unless you specify the profile. You can attach access profiles globally at the **[edit]** hierarchy level, or you can apply them to DHCP clients or subscribers, VLANs, or to a routing instance.

Options

profile-name—Name of the access profile that you configured at the **[edit access profile name]** hierarchy level.

Required Privilege Level

access—To view this statement in the configuration.

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

RELATED DOCUMENTATION

RADIUS Servers and Parameters for Subscriber Access

Attaching Access Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces

[Configuring Access Components for the DHCP Layer 3 Wholesale Network Solution | 16](#)

[Configuring Access Components for the PPPoE Wholesale Network Solution | 72](#)

access-profile (Dynamic VLAN)

Syntax

```
access-profile vlan-access-profile-name;
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges dynamic-profile profile-name]
```

Release Information

Statement introduced in Junos OS Release 16.2.

Description

Access profiles contain subscriber access authentication, authorization and accounting (AAA) configuration parameters. You can create an access profiles and then attach it at various configuration levels. When you attach an access profile to an interface configured for dynamic VLAN or stacked VLAN, all the VLANs and stacked VLANs use the same set of AAA parameters configured in that access profile. The different access profiles can have different authentication/authorization settings so you can, for example, have authentication on some VLAN or stacked VLAN ranges but no authentication on other ranges.

You can assign different access profiles to different dynamic profiles on the same interface. If you assign an access profile at the global level, but a different access profile is assigned at the interface level, the access profile at the interface level authenticates all dynamic VLANs and stacked VLANs on the interface. Access profiles can be assigned at various levels, but the most specific access profile takes precedence over any other profile assignments.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| [Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs](#)

access-profile (Dynamic Stacked VLAN)

Syntax

```
access-profile svlan-access-profile-name;
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure stacked-vlan-ranges dynamic-profile profile-name]
```

Release Information

Statement introduced in Junos OS Release 16.2.

Description

Access profiles contain subscriber access authentication, authorization and accounting (AAA) configuration parameters. You can create an access profiles and then attach it at various configuration levels. When you attach an access profile to an interface configured for dynamic VLAN or stacked VLAN, all the VLANs and stacked VLANs use the same set of AAA parameters configured in that access profile. The different access profiles can have different authentication/authorization settings so you can, for example, have authentication on some VLAN and stacked VLAN ranges but no authentication on other ranges.

You can assign different access profiles to different dynamic profiles on the same interface. If you assign an access profile at the global level, but a different access profile is assigned at the interface level, the access profile at the interface level authenticates all dynamic VLANs and stacked VLANs on the interface. Access profiles can be assigned at various levels, but the most specific access profile takes precedence over any other profile assignments..

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| [Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs](#)

active-server-group

Syntax

```
active-server-group server-group-name <allow-server-change>;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit forwarding-options dhcp-relay group group-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6 group group-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay
  dhcpv6],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay
  dhcpv6 group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay
  group group-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay]
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6 group group-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name]
```

Release Information

Statement introduced in Junos OS Release 8.3.

Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.

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

allow-server-change option added in Junos OS Release 16.2R1.

Description

Apply a DHCP relay agent configuration to the named group of DHCP server addresses. The server group itself is configured with the **server-group** statement. You can apply an active server group globally or for specific groups of interfaces, configured with the **group** statement. An active server group applied to an interface group overrides a global configuration.

Options

allow-server-change—(Optional) (DHCPv4 only) Enable the relay agent to accept and forward a DHCP request (renew or rebind) ACK message to the client from any DHCP local server in the active server

group. Starting in Junos OS Release 18.4R1, this option also applies to DHCP information request (DHCPINFORM) ACK messages.

Default: Forward ACK messages from only the original binding server.

server-group-name—Name of the group of DHCP or DHCPv6 server addresses to which the DHCP or DHCPv6 relay agent configuration applies.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview

Configuring Active Server Groups to Apply a Common DHCP Relay Agent Configuration to Named Server Groups

Configuring Group-Specific DHCP Relay Options

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 dlcidlcid-identifier;
  multipoint-destination address {
    epd-threshold cells;
    inverse-arp;
    oam-liveness {
      up-count cells;
      down-count cells;
    }
    oam-period (disable | seconds);
    shaping {
      (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate burst length);
      queue-length number;
    }
    vci vpi-identifier.vci-identifier;
  }
  primary;
  preferred;
  virtual-gateway-address
  (vrrp-group | vrrp-inet6-group) group-number {
    (accept-data | no-accept-data);
    advertise-interval seconds;
    authentication-type authentication;
    authentication-key key;
    fast-interval milliseconds;
    (preempt | no-preempt) {
      hold-time seconds;
    }
  }
  priority-number number;
  track {
    priority-cost seconds;
    priority-hold-time interface-name {
      interface priority;
      bandwidth-threshold bits-per-second {
        priority;

```

```
    }  
  }  
  route ip-address/mask routing-instance instance-name priority-cost cost;  
}  
virtual-address [ addresses ];  
}  
}
```

Hierarchy Level

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

NOTE: If you configure the same address on multiple interfaces in the same routing instance, Junos OS uses only the first configuration, and the remaining address configurations are ignored and can leave interfaces without an address. Interfaces that do not have an assigned address cannot be used as a donor interface for an unnumbered Ethernet interface.

For example, in the following configuration the address configuration of interface xe-0/0/1.0 is ignored:

```
interfaces {
  xe-0/0/0 {
    unit 0 {
      family inet {
        address 192.168.1.1/8;
      }
    }
  }
  xe-0/0/1 {
    unit 0 {
      family inet {
        address 192.168.1.1/8;
      }
    }
  }
}
```

For more information on configuring the same address on multiple interfaces, see *Configuring the Interface Address*.

- In Junos OS Release 13.3 and later, when you configure an IPv6 host address and an IPv6 subnet address on an interface, the commit operation fails.
- In releases earlier than Junos OS Release 13.3, when you use the same configuration on an interface, the commit operation succeeds, but only one of the IPv6 addresses that was entered is assigned to the interface. The other address is not applied.

Options

address—Address of the interface.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

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

<i>Configuring the Protocol Family</i>
<i>family</i>
<i>negotiate-address</i>
unnumbered-address (Ethernet) 486

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 {
    active-drain;
    family family {
      dhcp-attributes {
        protocol-specific attributes;
      }
      excluded-address ip-address;
      excluded-range name low minimum-value high maximum-value;
      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;
      }
    }
    hold-down;
    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: Support for subordinate statements is platform-specific. See individual statement topics for support information.

Options

abated-utilization—(ACX Series, MX Series only) Starting in Junos OS Release 11.2, generate SNMP traps for DHCP address pools or linked set of address pools. No SNMP traps are generated unless a value is configured. Default: Abated utilization is not set. Delete the abated-utilization value to unset.

Values: *percentage*—Threshold below which an SNMP trap clear is generated. Range: 1 through 98.

abated-utilization-v6—(ACX Series, MX Series only) Starting in Junos OS Release 11.2, generate SNMP traps for DHCPv6 address pools or linked set of address pools. No SNMP traps are generated unless a value is configured. Default: Abated utilization is not set. Delete the abated-utilization value to unset.

Values: *percentage*—Threshold below which an SNMP trap clear is generated. Range: 1 through 98.

high-utilization—(ACX Series, MX Series only) Starting in Junos OS Release 11.2, generate an SNMP trap when the DHCP address pool or linked set of address pools use surpasses the specified percentage. Default: High utilization is not set. Delete the high-utilization value to unset.

Values: *percentage*—Percentage used to generate a trap. Range: 2 through 99.

high-utilization-v6—(ACX Series, MX Series only) Starting in Junos OS Release 11.2, generate an SNMP trap when the DHCPv6 address pool or linked set of address pools use surpasses the specified percentage. Default: High utilization is not set. Delete the high-utilization value to unset.

Values: *percentage*—Percentage used to generate a trap. Range: 2 through 99.

neighbor-discovery-router-advertisement—(M Series, MX Series, SRX Series, T Series only) Configure the name of the address-assignment pool used to assign the router advertisement prefix.

Values: *ndra-pool-name*—Name of the address-assignment pool.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Address-Assignment Pools Overview

Address-Assignment Pool Configuration Overview

Configuring an Address-Assignment Pool for L2TP LNS with Inline Services

Configuring Address-Assignment Pool Usage Threshold Traps

Configuring an Address-Assignment Pool Used for Router Advertisements

adjacency-loss-hold-time (ANCP)

Syntax

```
adjacency-loss-hold-time seconds;
```

Hierarchy Level

```
[edit protocols ancp],  
[edit protocols ancp neighbor ip-address]
```

Release Information

Statement introduced in Junos OS Release 16.1R4.

Description

Configure the ANCP agent to monitor, either globally or for a specified neighbor, how long an ANCP adjacency is down and to trigger a state change for the subscriber access line if the hold timer expires before the adjacency comes back up, as indicated by a Port Up message on the access line. By default, there is no delay between detecting an adjacency loss and triggering the state change.

Options

seconds—Duration of period that the ANCP agent monitors loss of adjacency.

Default: 0 seconds

Range: 0 through 1800 seconds

Required Privilege Level

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring the ANCP Agent to Dampen the Effects of Short-Term Adjacency Losses](#) | 144

[Configuring ANCP Neighbors](#)

[Configuring the ANCP Agent](#)

[Layer 2 Wholesale with ANCP-Triggered VLANs Overview](#) | 116

ancp

Syntax

```

ancp {
  adjacency-loss-hold-time seconds;
  adjacency-timer seconds;
  gsmp-syn-timeout seconds;
  gsmp-syn-wait;
  interfaces {
    interface-set interface-set-name {
      access-identifier identifier-string;
      underlying-interface underlying-interface-name;
    }
    interface-name {
      access-identifier identifier-string;
    }
  }
  maximum-discovery-table-entries entry-number;
  maximum-helper-restart-time;
  neighbor ip-address {
    adjacency-loss-hold-time seconds;
    adjacency-timer;
    auto-configure-trigger interface interface-name;
    ietf-mode;
    maximum-discovery-table-entries entry-number;
    pre-ietf-mode;
  }
  pre-ietf-mode;
  qos-adjust {
    adsl-bytes bytes;
    adsl2-bytes bytes;
    adsl2-plus-bytes bytes;
    other-bytes bytes;
    other-overhead-adjust percentage;
    sdsl-bytes bytes;
    sdsl-overhead-adjust percentage;
    vdsl-bytes bytes;
    vdsl-overhead-adjust percentage;
    vdsl2-bytes bytes;
    vdsl2-overhead-adjust percentage;
  }
  qos-adjust-adsl adjustment-factor;
  qos-adjust-adsl2 adjustment-factor;
  qos-adjust-adsl2-plus adjustment-factor;

```

```

qos-adjust-other adjustment-factor;
qos-adjust-sdsl adjustment-factor;
qos-adjust-vdsl adjustment-factor;
qos-adjust-vdsl2 adjustment-factor;
traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
}
}

```

Hierarchy Level

[edit protocols]

Release Information

Statement introduced in Junos OS Release 9.4.

Description

Configure Junos OS ANCP agent features.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

NOTE: When you deactivate the ANCP protocol, the router does not perform a commit check to determine whether any ANCP or L2-BSA subscribers are present (active or inactive). Any subscribers that are active at the time of deactivation remain active.

Required Privilege Level

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| *Configuring the ANCP Agent*

authentication

Syntax

```
authentication {
  packet-types [packet-types];
  password password-string;
  username-include {
    circuit-id;
    circuit-type;
    delimiter delimiter-character;
    domain-name domain-name-string;
    interface-name;
    mac-address;
    option-18;
    option-37;
    option-82 <circuit-id> <remote-id>;
    radius-realm radius-realm-string;
    remote-id;
    user-prefix user-prefix-string;
    vlan-tags;
  }
}
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges],
[edit interfaces interface-name auto-configure stacked-vlan-ranges]
```

Release Information

Statement introduced in Junos OS Release 10.0.

Description

Specify the authentication parameters that trigger the Access-Request message to AAA for the interface.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Subscribers over Static Interfaces Configuration Overview

Configuring the Static Subscriber Global Authentication Password

[Configuring a Username for Authentication of Out-of-Band Triggered Dynamic VLANs | 139](#)

[Layer 2 Wholesale with ANCP-Triggered VLANs Overview | 116](#)

authentication (DHCP Local Server)

Syntax

```
authentication {
  password password-string;
  username-include {
    circuit-type;
    client-id;
    delimiter delimiter-character;
    domain-name domain-name-string;
    interface-description (device-interface | logical-interface);
    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;
    vlan-tags;
  }
}
```

Hierarchy Level

```
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dual-stack-group dual-stack-group-name],
[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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| *Specifying Authentication Support*

authentication (DHCP Relay Agent)

Syntax

```
authentication {
  password password-string;
  username-include {
    circuit-type;
    client-id;
    delimiter delimiter-character;
    domain-name domain-name-string;
    interface-description (device-interface | logical-interface);
    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;
    vlan-tags;
  }
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name],
[edit forwarding-options dhcp-relay group group-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Release Information

Statement introduced in Junos OS Release 9.1.

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

Support at the **[edit ... dhcpv6]** hierarchy levels introduced in Junos OS Release 11.4.

Support at the **[edit ... dual-stack-group dual-stack-group-name]** hierarchy level introduced in Junos OS Release 15.1.

Description

Configure the parameters the router sends to the external AAA server. A group configuration takes precedence over a global DHCP relay configuration. Use the statement at the [\[edit...dhcpv6\]](#) hierarchy levels to configure DHCPv6 support.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[dhcp-relay | 264](#)

Specifying Authentication Support

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 added in Junos OS Release 11.2.

nasreq option added in Junos OS Release 16.1.

s6a option added in Junos OS Release 19.3R1.

Description

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

A given subscriber does not undergo both authentication and authorization as separate steps. When both **authentication-order** and **authorization-order** are specified, DHCP subscribers honor the configured authorization order, all other subscribers use the configured authentication-order.

Starting in Junos OS Release 18.2R1, the **password** option can also be used to specify that local authentication and local authorization is attempted for individual subscribers that are configured with the **subscriber** statement at the `[edit access profile profile-name]` hierarchy level.

Options

authentication-methods—Ordered list of methods to use for authentication attempts. The list includes one or more of the following methods in any combination:

- **nasreq**—Verify subscribers using the Diameter-based Network Access Server Requirements (NASREQ) protocol.
- **none**—No authentication is performed. 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.

NOTE: Subscriber access management does not support the **none** option; authentication fails when this option is specified.

- **password**—Verify the client using the information configured at the [edit access profile *profile-name* client *client-name*] hierarchy level.

Subscriber access management does not support the **password** option until Junos OS Release 18.2R1. Starting in Junos OS Release 18.2R1, this option is used to enable local authentication and optionally local authorization for individual subscribers. Local authentication is typically used when you do not have external authentication and authorization servers. The password itself must be configured with the **subscriber** statement in the same access profile. Local authentication is performed when a subscriber logs in with a matching username; it succeeds if the subscribers login password matches the password in the profile.

If you have external authentication and authorization servers, you can use local authentication as a backup authentication method. In this case, configure **password** other than first in the list of methods.

- **radius**—Verify the client using RADIUS authentication services.
- **s6a**—Verify subscribers using the Diameter-based s6a protocol.

Default: password

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

Configuring Local Authentication and Authorization for Subscribers

Example: Configure S6a Application

auto-configure

Syntax

```

auto-configure {
  vlan-ranges {
    access-profile profile-name;
    authentication {
      packet-types [packet-types];
      password password-string;
      username-include{
        circuit-id;
        circuit-type;
        delimiter delimiter-character;
        domain-name domain-name-string;
        interface-name;
        mac-address;
        option-18;
        option-37;
        option-82 <circuit-id> <remote-id>;
        radius-realm radius-realm-string;
        remote-id;
        user-prefix user-prefix-string;
        vlan-tags;
      }
    }
    dynamic-profile profile-name {
      accept (any | dhcp-v4 | dhcp-v6 | inet | inet6 | pppoe);
      accept-out-of-band protocol;
      ranges (any | low-tag)-(any | high-tag);
    }
    override;
  }
  stacked-vlan-ranges {
    access-profile profile-name;
    authentication {
      packet-types [packet-types];
      password password-string;
      username-include {
        circuit-type;
        delimiter delimiter-character;
        domain-name domain-name-string;
        interface-name;
        mac-address;
        option-18;

```

```

    option-37;
    option-82 <circuit-id> <remote-id>;
    radius-realm radius-realm-string;
    user-prefix user-prefix-string;
    vlan-tags;
  }
}
dynamic-profile profile-name {
  accept (any | dhcp-v4 | dhcp-v6 | inet | inet6 | pppoe);
  ranges (any | low-tag-high-tag),(any | low-tag-high-tag);
}
override;
}
remove-when-no-subscribers;
}

```

Hierarchy Level

```
[edit interfaces interface-name]
```

Release Information

Statement introduced in Junos OS Release 9.5.

Description

Enable the configuration of dynamic, auto-sensed VLANs.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs](#)

[Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs](#)

auto-configure-trigger interface (ANCP)

Syntax

```
auto-configure-trigger interface interface-name;
```

Hierarchy Level

```
[edit protocols ancps neighbor ip-address]
```

Release Information

Statement introduced in Junos OS Release 16.1R4.

ps interface type support added in Junos OS Release 19.3R1.

Description

Map an ANCP neighbor to a subscriber-facing physical interface on the router, so that ANCP Port Up and Port Down messages trigger notifications to the auto-configuration daemon (autoconfd) to initiate VLAN creation (Port Up) or removal (Port Down).

Options

interface-name—Name of the physical interface. Can be any of the following interface types: ae, ge, xe, et, demux, ps.

Required Privilege Level

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring ANCP Neighbors

Configuring the ANCP Agent

backup-on-failure (Accounting Options)

Syntax

```
backup-on-failure (master-and-slave | master-only);
```

Hierarchy Level

[edit accounting-options *file filename*]

Release Information

Statement introduced in Junos OS Release 16.1.

Description

Configure the router to save a copy of the accounting file locally, to the `/var/log/pfedBackup` directory of the relevant Routing Engine, in the event that file transfer to the remote archive sites cannot be completed.

Options

master-and-slave—Back up accounting files from both the master Routing Engine and the backup Routing Engine.

master-only—Back up accounting files from only the master Routing Engine.

Required Privilege Level

system—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Accounting-Data Log Files

[Configuring Flat-File Accounting for Layer 2 Wholesale | 153](#)

[Configuring Flat-File Accounting for Extensible Subscriber Services Management | 158](#)

[Flat-File Accounting Overview | 149](#)

circuit-id (VLAN Authentication Username)

Syntax

```
circuit-id;
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges authentication username-include]
```

Release Information

Statement introduced in Junos OS Release 16.1R4.

Description

Include the agent circuit identifier (ACI) in the username sent to RADIUS for authentication of the dynamic VLAN. The ACI is conveyed by the Access-Loop-Circuit-ID TLV in an out-of-band ANCP Port Up message.

NOTE: This statement is not supported for stacked VLANs.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring a Username for Authentication of Out-of-Band Triggered Dynamic VLANs | 139](#)

[Configuring VLAN Interface Username Information for AAA Authentication](#)

[Layer 2 Wholesale with ANCP-Triggered VLANs Overview | 116](#)

cleanup-interval (Accounting Options)

Syntax

```
cleanup-interval days;
```

Hierarchy Level

[edit accounting-options]

Release Information

Statement introduced in Junos OS Release 16.1.

Description

Configure the interval to delete files from the local backup directory.

Options

days—Number of days after which accounting-options files are to be deleted from the backup directory.

Range: 1 through 31 days

Default: 1 day

Required Privilege Level

system—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Accounting-Data Log Files

[Configuring Flat-File Accounting for Layer 2 Wholesale | 153](#)

[Configuring Flat-File Accounting for Extensible Subscriber Services Management | 158](#)

[Flat-File Accounting Overview | 149](#)

compress (Accounting Options)

Syntax

```
compress;
```

Hierarchy Level

[edit accounting-options file *filename*]

Release Information

Statement introduced in Junos OS Release 16.1.

Description

Compress the accounting file during file transfer to the backup site.

Required Privilege Level

system—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Accounting-Data Log Files

[Configuring Flat-File Accounting for Layer 2 Wholesale | 153](#)

[Configuring Flat-File Accounting for Extensible Subscriber Services Management | 158](#)

[Flat-File Accounting Overview | 149](#)

connectivity-type

Syntax

```
connectivity-type (ce | irb | permanent);
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 9.1.

irb option introduced in Junos OS Release 9.3.

permanent option introduced in Junos OS Release 10.4.

Description

Specify when a VPLS connection is taken down depending on whether or not the interface for the VPLS routing instance is customer-facing or integrated routing and bridging (IRB).

NOTE: The **connectivity-type** statement is not supported for FEC 129 VPLS (also known as LDP VPLS with BGP-based autodiscovery).

Default

ce

Options

ce—Require that for the VPLS connection to be up, the customer-facing interface for the VPLS routing instance must also be up. If the customer-facing interface fails, the VPLS connection is taken down.

irb—Allow a VPLS connection to remain up so long as an IRB interface is configured for the VPLS routing instance.

permanent—Allow a VPLS connection to remain up until specifically taken down. This option is reserved for use in configuring Layer 2 Wholesale subscriber networks. See the *Broadband Subscriber Management Solutions Guide* for details about configuring a Layer 2 Wholesale network.

NOTE: To specifically take down a VPLS routing instance that is using the **permanent** option, all associated static logical interfaces must also be down.

Required Privilege Level

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring VPLS Routing Instance and VPLS Interface Connectivity

[Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers | 105](#)

core-facing**Syntax**

```
core-facing;
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 11.2.

Description

Specifies that the VLAN is physically connected to a core-facing ISP router and ensures that the network does not improperly treat the interface as a client interface. When specified, the interface is inserted into the core-facing default mesh group where traffic from pseudowires that belong to the default mesh group is not forwarded on the core-facing link.

Required Privilege Level

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Direct ISP-Facing Interfaces for the Layer 2 Wholesale Solution | 101](#)

demux0 (Dynamic Interface)

Syntax

```

demux0 {
  unit logical-unit-number {
    demux-options {
      underlying-interface interface-name
    }
    family family {
      access-concentrator name;
      address address;
      demux-source {
        source-prefix;
      }
      direct-connect;
      duplicate-protection;
      dynamic-profile profile-name;
      filter {
        input filter-name;
        output filter-name;
      }
      mac-validate (loose | strict);
      max-sessions number;
      max-sessions-vs-a-ignore;
      rpf-check {
        fail-filter filter-name;
        mode loose;
      }
      service-name-table table-name
      short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max maximum-seconds>;
      unnumbered-address interface-name <preferred-source-address address>;
    }
    filter {
      input filter-name;
      output filter-name;
    }
    vlan-id number;
  }
}

```

Hierarchy Level

[edit [dynamic-profiles](#) profile-name [interfaces](#)]

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Configure the logical demultiplexing (demux) interface in a dynamic profile.

Logical IP demux interfaces do not support IPv4 and IPv6 dual stack.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles

Demultiplexing Interface Overview

demux-options (Dynamic Interface)

Syntax

```
demux-options {  
    underlying-interface interface-name  
}
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Configure logical demultiplexing (demux) interface options in a dynamic profile.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles

Demultiplexing Interface Overview

demux-source (Dynamic IP Demux Interface)

Syntax

```
demux-source {  
    source-address;  
}
```

Hierarchy Level

[edit [dynamic-profiles](#) *profile-name* [interfaces demux0](#) *unit* *logical-unit-number* [family](#) *family*]

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Configure a logical demultiplexing (demux) source address for a subscriber in a dynamic profile.

Options

source-address—Either the specific source address you want to assign to the subscriber interface or the source address variable. For IPv4, specify **\$junos-subscriber-ip-address**; for IPv6, specify **\$junos-subscriber-ipv6-address**. The source address for the interface is dynamically supplied by DHCP when the subscriber accesses the router.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles

Demultiplexing Interface Overview

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.

demux-source (Underlying Interface)

Syntax

```
demux-source family;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name routing-instances routing-instance-name interfaces interface-name unit
logical-unit-number]
```

Release Information

Statement introduced in Junos OS Release 9.0.

Support for aggregated Ethernet added in Junos OS Release 9.4.

Description

Configure the logical demultiplexing (demux) source family type on the IP demux underlying interface.

NOTE: The IP demux interface feature currently supports only Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet underlying interfaces.

Options

family—Protocol family:

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

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring an IP Demultiplexing Interface

Configuring a VLAN Demultiplexing Interface

dhcp-attributes (Address-Assignment Pools)

Syntax

```
dhcp-attributes {
  boot-file filename;
  boot-server (address | hostname);
  dns-server [ ipv6-address ];
  domain-name domain-name;
  exclude-prefix-len exclude-prefix-length;
  grace-period seconds;
  maximum-lease-time seconds;
  name-server [ server-list ];
  netbios-node-type node-type;
  option {
    [ (id-number option-type option-value)
      (id-number array option-type option-value) ];
  }
  option-match {
    option-82 {
      circuit-id value range named-range;
      remote-id value range named-range;
    }
  }
  preferred-lifetime seconds;
  router [ router-address ];
  server-identifier ip4-address;
  sip-server-address [ ipv6-address ];
  sip-server-domain-name domain-name;
  t1-percentage percentage;
  t1-renewal-time;
  t2-percentage percentage;
  t2-rebinding-time;
  tftp-server address;
  valid-lifetime seconds;
  wins-server [ servers ];
}
```

Hierarchy Level

```
[edit access address-assignment pool pool-name family family]
```

Release Information

Statement introduced in Junos OS Release 9.0.

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

exclude-prefix-len statement introduced in Junos OS Release 17.3 for MX Series.

Description

Configure DHCP attributes for the protocol family in a specific address pool. The attributes determine options and behaviors for the DHCP clients.

Options

boot-file—(EX Series, M Series, MX Series, SRX Series, T Series only) Set the boot file advertised to DHCP clients. After the client receives an IP address and the boot file location from the DHCP server, the client uses the boot image stored in the boot file to complete DHCP setup. This configuration is equivalent to DHCP option 67.

Values: *filename*—Location of the boot file on the boot server. The filename can include a pathname.

boot-server—(EX Series, M Series, MX Series, SRX Series, T Series only) Configure the name of the boot server advertised to DHCP clients. The client uses a boot file located on the boot server to complete DHCP setup. This configuration is equivalent to DHCP option 66.

Values:

- *address*—IPv4 address of a boot server.
- *hostname*—Fully qualified hostname of a boot server.

dns-server—(MX Series only) Specify a DNS server to which clients can send DNS queries. This is equivalent to DHCPv6 option 23. To specify multiple DNS servers, add multiple **dns-server** statements in order of preference.

Values: *ipv6-address*—IPv6 address of a DNS server.

domain-name—(EX Series, MX Series only) Configure the name of the domain in which clients search for a DHCP server host. This is the default domain name that is appended to hostnames that are not fully qualified. This is equivalent to DHCP option 15.

Values: *domain-name*—Name of the domain.

exclude-prefix-len *exclude-prefix-length*—Specify the length of the IPv6 prefix to be excluded from the delegated prefix. Range: 1 through 128.

grace-period—(M Series, MX Series, SRX Series, T Series only) 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.

Values: *seconds*—Number of seconds the lease is retained.

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

Default: 0 (no grace period).

maximum-lease-time—(EX Series, M120, MX Series, SRX240, SRX3400, T640, T1600 only) 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. The **maximum-lease-time** is mutually exclusive with both the **preferred-lifetime** and the **valid-lifetime**, and cannot be configured with either timer.

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

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

Default: 86,400 (24 hours).

name-server—(M120, MX Series, SRX Series, T640, T1600 only) Configure one or more Domain Name System (DNS) name servers available to the client to resolve hostname-to-client mappings. This is equivalent to DHCP option 6.

Values: *server-names*—IP addresses of the domain name servers, listed in order of preference.

netbios-node-type—(M Series, MX Series, SRX Series, T Series only) Specify the NetBIOS node type. This is equivalent to DHCP option 46.

Values: *node-type*—One of the following node types:

- *b-node*—Broadcast node.
- *h-node*—Hybrid node.
- *m-node*—Mixed node.
- *p-node*—Peer-to-peer node.

option—(EX Series, M Series, MX Series, SRX Series, T Series only) Specify user-defined options that are added to client packets. Starting in Junos OS Release 13.3, the **hex-string** option type was introduced.

Values:

- *array*—An option can include an array of option types.
- *id-number*—Any whole number. The ID number is used to index the option and must be unique across a DHCP server.
- *option-type*—Any of the following types: byte, byte-stream, flag, hex-string, integer, ip-address, short, string, unsigned-integer, or unsigned-short.
- *option-value*—Value associated with an option. The option value must be compatible with the option type (for example, an On or Off value for a flag type).

preferred-lifetime—(EX Series, M Series, MX Series only) Starting in Junos OS Release 13.3, specify the length of time, in seconds, that the DHCPv6 server keeps the IPv6 prefix active. When the lifetime expires, the address is deprecated. If the **valid-lifetime** is also configured, the **preferred-lifetime** must be less than the **valid-lifetime**. The **preferred-lifetime** and the **maximum-lease-time** are mutually exclusive and cannot both be configured.

Values: *seconds*—Number of seconds that the IPv6 prefix is active.

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

Default: 86,400 (24 hours).

router—(EX Series, MX Series only) Specify one or more routers located on the client's subnet. This statement is the equivalent of DHCP option 3.

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

server-identifier—(EX Series, MX Series, SRX Series only) Specify the IP address that is used as the source address the DHCP server includes in IP packets when communicating with clients. The address is included in the DHCP packet in option 54.

Values: *ipv4-address*—IP address.

sip-server-address—(EX Series, M Series, MX Series, SRX Series, T Series only) Specify a SIP outbound proxy server that DHCPv6 local server clients can use. This is equivalent to DHCPv6 option 22. To specify multiple servers, add multiple **sip-server-address** statements in order of preference.

Values: *ipv6-address*—IPv6 address of a SIP outbound proxy server.

sip-server-domain-name—(M Series, MX Series, SRX Series, T Series only) Configure the domain name of the SIP outbound proxy server that DHCPv6 local server clients can use. This is equivalent to DHCPv6 option 21.

Values: *domain-name*—Name of the domain.

t1-percentage—(EX Series, M Series, MX Series only) Specify a percentage of the **preferred-lifetime** value. After this percentage of the **preferred-lifetime** value elapses, the DHCPv4 or DHCPv6 client requests an extension on its lease from the originating DHCP local server. The **t1-percentage** is also referred to as the renewal time. The **t1-percentage** value must be less than the **t2-percentage** value. DHCPv4 server support was added in Junos OS Release 17.2.

Values: *percentage*—Percentage of the **preferred-lifetime** value.

Range: 0 through 100.

Default: If the **t1-percentage** value is not configured, the default is based on the **preferred-lifetime** value:

- If the **preferred-lifetime** value is finite, the default is 50 percent of the **preferred-lifetime** value.
- If the **preferred-lifetime** value is infinite, the default is also infinite.

t1-renewal-time—(MX Series only) Starting in Junos OS Release 17.2, specify the time (T1) at which the DHCPv4 or DHCPv6 client requests an extension (renewal) of the existing lease. This time is expressed as the number of seconds since the beginning of the lease. Using this statement to configure a duration in seconds is an alternative to using the **t1-percentage** statement.

Values: *seconds*—Number of seconds.

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

Default: 50 percent of the lease duration (**preferred-lifetime**).

t2-percentage—(EX Series, M Series, MX Series only) Starting in Junos OS Release 13.3, specify a percentage of the **preferred-lifetime** value. After this percentage of the **preferred-lifetime** value elapses, the DHCPv4 or DHCPv6 client requests an extension on its lease from any available DHCPv4 or DHCPv6 server. The **t2-percentage** is also referred to as the rebinding time. The **t2-percentage** value must be greater than the **t1-percentage** value. DHCPv4 server support was added in Junos OS Release 17.2.

Values: *percentage*—Percentage of the **preferred-lifetime** value.

Range: 0 through 100.

Default: Default: If the **t2-percentage** value is not configured, the default is based on the **preferred-lifetime** value:

- If the **preferred-lifetime** value is finite, the default is 80 percent of the **preferred-lifetime** value.
- When the **preferred-lifetime** value is infinite, the default is also infinite.

t2-rebinding-time—(MX Series only) Starting in Junos OS Release 17.2, specify the time (T2) at which the DHCPv4 or DHCPv6 client attempts to contact any DHCP server to request an extension (rebinding) of the existing lease. This time is expressed as the number of seconds since the beginning of the lease. Using this statement to configure a duration in seconds is an alternative to using the **t2-percentage** statement.

Values: *seconds*—Number of seconds.

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

Default: The default value depends on the client:

- (DHCPv4 clients) 87.5 percent of the lease duration (**preferred-lifetime**).
- (DHCPv6 clients) 80 percent of the lease duration (**preferred-lifetime**).

tftp-server—(ACX Series, SRX Series only) Specify the Trivial File Transfer Protocol (TFTP) server that the client uses to obtain the client configuration file. This is equivalent to DHCP option 150.

Values: *ip-address*—IP address of the TFTP server.

valid-lifetime—(EX Series, M Series, MX Series only) Starting in Junos OS Release 13.3, specify the length of time, in seconds, that the DHCPv6 server keeps the IPv6 prefix valid. When the lifetime expires, the address becomes invalid. If the **preferred-lifetime** is also configured, the **valid-lifetime** must be greater than the **preferred-lifetime**. The **valid-lifetime** and the **maximum-lease-time** are mutually exclusive and cannot both be configured.

Values: *seconds*—Number of seconds that the IPv6 prefix is valid.

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

Default: 86,400 (24 hours).

wins-server—(M Series, MX Series, SRX Series, T Series only) Specify one or more NetBIOS name servers (NBNS) that the client uses to resolve NetBIOS names. This is equivalent to DHCP option 44.

Values: *ipv4-address*—IP address of each NetBIOS name server. Add them to the configuration in order of preference.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Address-Assignment Pools Overview

Attributes That Can Be Applied to DHCP Clients

Address-Assignment Pool Configuration Overview

Configuring DHCP Client-Specific Attributes Applied When Clients Obtain an Address

DHCP Lease Timers

Configuring DHCP Attributes for All Clients or a Group of Clients

dhcp-local-server

Syntax

```

dhcp-local-server {
  access-profile profile-name;
  allow-active-leasequery {
    idle-timeout seconds;
    peer-address address;
    timeout seconds;
  }
  allow-bulk-leasequery {
    max-connections number-of-connections;
    max-empty-replies number-of-replies;
    restricted-requestor;
    timeout seconds;
  }
  allow-leasequery {
    restricted-requestor;
  }
  authentication {
    password password-string;
    username-include {
      circuit-type;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-description (device-interface | logical-interface);
      interface-name ;
      logical-system-name;
      mac-address;
      option-60;
      option-82 <circuit-id> <remote-id>;
      routing-instance-name;
      user-prefix user-prefix-string;
      vlan-tags;
    }
  }
  dhcpv6 {
    access-profile profile-name;
    allow-active-leasequery {
      idle-timeout seconds;
      peer-address address;
      timeout seconds;
    }
    allow-bulk-leasequery {

```

```
max-connections number-of-connections;  
max-empty-replies number-of-replies;  
restricted-requestor;  
timeout seconds;  
}  
allow-leasequery {  
    restricted-requestor;  
}  
authentication {  
    ...  
}  
duplicate-clients incoming-interface;
```

```

group group-name {
  access-profile profile-name;
  authentication {
    ...
  }
  interface interface-name {
    access-profile profile-name;
    exclude;
    overrides {
      asymmetric-lease-time seconds;
      asymmetric-prefix-lease-time seconds;
      delay-advertise {
        based-on (option-15 | option-16 | option-18 | option-37) {
          equals {
            ascii ascii-string;
            hexadecimal hexadecimal-string;
          }
          not-equals {
            ascii ascii-string;
            hexadecimal hexadecimal-string;
          }
          starts-with {
            ascii ascii-string;
            hexadecimal hexadecimal-string;
          }
        }
        delay-time seconds;
      }
      dual-stack dual-stack-group-name;
      interface-client-limit number;
      multi-address-embedded-option-response;
      process-inform {
        pool pool-name;
      }
      protocol-attributes attribute-set-name;
      rapid-commit;
    }
    service-profile dynamic-profile-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    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;  
  }  
  layer2-liveness-detection {  
    max-consecutive-retries number;  
    transmit-interval interval;  
  }  
}
```

```

overrides {
  asymmetric-lease-time seconds;
  asymmetric-prefix-lease-time seconds;
  delay-advertise {
    based-on (option-15 | option-16 | option-18 | option-37) {
      equals {
        ascii ascii-string;
        hexadecimal hexadecimal-string;
      }
      not-equals {
        ascii ascii-string;
        hexadecimal hexadecimal-string;
      }
      starts-with {
        ascii ascii-string;
        hexadecimal hexadecimal-string;
      }
    }
    delay-time seconds;
  }
  delegated-pool;
  dual-stack dual-stack-group-name;
  interface-client-limit number;
  multi-address-embedded-option-response;
  process-inform {
    pool pool-name;
  }
  protocol-attributes attribute-set-name;
  rapid-commit;
}
route-suppression;
server-duid-type type;
service-profile dynamic-profile-name;
short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}

```

```
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;  
    }  
    layer2-liveness-detection {  
      max-consecutive-retries number;  
      transmit-interval interval;  
    }  
  }  
}
```



```

overrides {
  asymmetric-lease-time seconds;
  asymmetric-prefix-lease-time seconds;
  delay-advertise {
    based-on (option-15 | option-16 | option-18 | option-37) {
      equals {
        ascii ascii-string;
        hexadecimal hexadecimal-string;
      }
      not-equals {
        ascii ascii-string;
        hexadecimal hexadecimal-string;
      }
      starts-with {
        ascii ascii-string;
        hexadecimal hexadecimal-string;
      }
    }
    delay-time seconds;
  }
  delegated-pool;
  dual-stack dual-stack-group-name;
  include-option-82 {
    forcerenew;
    nak;
  }
  interface-client-limit number;
  multi-address-embedded-option-response;
  process-inform {
    pool pool-name;
  }
  protocol-attributes attribute-set-name;
  rapid-commit;
}

```

```
reconfigure {  
    attempts attempt-count;  
    clear-on-abort;  
    strict;  
    support-option-pd-exclude;  
    timeout timeout-value;  
    token token-value;  
    trigger {  
        radius-disconnect;  
    }  
}  
reauthenticate (<lease-renewal> <remote-id-mismatch >);  
requested-ip-network-match subnet-mask;  
route-suppression;  
server-duid-type type;  
service-profile dynamic-profile-name;  
short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;  
}
```

```

dual-stack-group name {
  access-profile access-profile;
  authentication {
    password password-string;
    username-include {
      circuit-type;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-description (device-interface | logical-interface);
      interface-name ;
      logical-system-name;
      mac-address;
      relay-agent-interface-id;
      relay-agent-remote-id;
      routing-instance-name;
      user-prefix user-prefix-string;
      vlan-tags;
    }
  }
  classification-key {
    circuit-id circuit-id;
    mac-address mac-address;
    remote-id remote-id;
  }
  dual-stack-interface-client-limit number;
  dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
  }
  liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
      layer2-liveness-detection {
        max-consecutive-retries number;
        transmit-interval interval;
      }
    }
  }
  on-demand-address-allocation;
  protocol-master (inet | inet6);
  reauthenticate (<lease-renewal> <remote-id-mismatch >);
  service-profile service-profile;
  short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}

```

```
duplicate-clients-in-subnet (incoming-interface | option-82);  
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;
    overrides {
      asymmetric-lease-time seconds;
      client-discover-match (option60-and-option82 | incoming-interface);
      delay-offer {
        based-on (option-60 | option-77 | option-82) {
          equals {
            ascii ascii-string;
            hexadecimal hexadecimal-string;
          }
          not-equals {
            ascii ascii-string;
            hexadecimal hexadecimal-string;
          }
          starts-with {
            ascii ascii-string;
            hexadecimal hexadecimal-string;
          }
        }
        delay-time seconds;
      }
      include-option-82 {
        forcerenew;
        nak;
      }
      dual-stack dual-stack-group-name;
      interface-client-limit number;
      process-inform {
        pool pool-name;
      }
      protocol-attributes attribute-set-name;
    }
    service-profile dynamic-profile-name;
    short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
    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;
  }
  layer2-liveness-detection {
    max-consecutive-retries number;
    transmit-interval interval;
  }
}
```

```

overrides {
  asymmetric-lease-time seconds;
  client-discover-match (option60-and-option82 | incoming-interface);
  delay-offer {
    based-on (option-60 | option-77 | option-82) {
      equals {
        ascii ascii-string;
        hexadecimal hexadecimal-string;
      }
      not-equals {
        ascii ascii-string;
        hexadecimal hexadecimal-string;
      }
      starts-with {
        ascii ascii-string;
        hexadecimal hexadecimal-string;
      }
    }
    delay-time seconds;
  }
  include-option-82 {
    forcerenew;
    nak;
  }
  dual-stack dual-stack-group-name;
  interface-client-limit number;
  process-inform {
    pool pool-name;
  }
  protocol-attributes attribute-set-name;
}
requested-ip-network-match subnet-mask
route-suppression;
service-profile dynamic-profile-name;
short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}

```

```

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;
    }
    layer2-liveness-detection {
      max-consecutive-retries number;
      transmit-interval interval;
    }
  }
}
on-demand-address-allocation;

```



```

overrides {
  asymmetric-lease-time seconds;
  client-discover-match <option60-and-option82 | incoming-interface>;
  delay-offer {
    based-on (option-60 | option-77 | option-82) {
      equals {
        ascii ascii-string;
        hexadecimal hexadecimal-string;
      }
      not-equals {
        ascii ascii-string;
        hexadecimal hexadecimal-string;
      }
      starts-with {
        ascii ascii-string;
        hexadecimal hexadecimal-string;
      }
    }
    delay-time seconds;
  }
  dual-stack dual-stack-group-name;
  interface-client-limit number;
  process-inform {
    pool pool-name;
  }
  protocol-attributes attribute-set-name;
}
pool-match-order {
  external-authority;
  ip-address-first;
  option-82;
}
protocol-master;
reauthenticate (<lease-renewal> <remote-id-mismatch >);
reconfigure {
  attempts attempt-count;
  clear-on-abort;
  strict;
  timeout timeout-value;
  token token-value;
  trigger {
    radius-disconnect;
  }
}

```

```

requested-ip-network-match subnet-mask;
route-suppression;
service-profile dynamic-profile-name;
short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}

```

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.

Statement introduced in Junos OS Release 13.2X51 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure Dynamic Host Configuration Protocol (DHCP) local server options on the router or switch to 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, therefore, 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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Understanding Differences Between Legacy DHCP and Extended DHCP

DHCPv6 Local Server Overview

dhcp-relay

Syntax

```

dhcp-relay {
  access-profile profile-name;
  active-leasequery {
    idle-timeout seconds;
    peer-address address;
    timeout seconds;
    topology-discovery;
  }
  active-server-group server-group-name;
  authentication {
    password password-string;
    username-include {
      circuit-type;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-description (device-interface | logical-interface);
      interface-name;
      logical-system-name;
      mac-address;
      option-60;
      option-82 <circuit-id> <remote-id>;
      routing-instance-name;
      user-prefix user-prefix-string;
      vlan-tags;
    }
  }
  bulk-leasequery {
    attempts number-of-attempts;
    timeout seconds;
  }
  dhcpv6 {
    access-profile profile-name;
    active-leasequery {
      idle-timeout seconds;
      peer-address address;
      timeout seconds;
      topology-discovery;
    }
    active-server-group server-group-name;
  }
  authentication {

```

```

password password-string;
username-include {
    circuit-type;
    client-id;
    delimiter delimiter-character;
    domain-name domain-name-string;
    interface-description (device-interface | logical-interface);
    interface-name interface-name;
    logical-system-name;
    mac-address mac-address;
    relay-agent-interface-id;
    relay-agent-remote-id;
    relay-agent-subscriber-id;
    routing-instance-name;
    user-prefix user-prefix-string;
    vlan-tags;
}
}
bulk-leasequery {
    attempts number-of-attempts;
    timeout seconds;
    trigger automatic;
}
duplicate-clients incoming-interface;
dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
}
forward-only {
    logical-system <current | default | logical-system-name>;
    routing-instance <current | default | routing-instance-name>;
}
forward-only-replies;
}
forward-snooped-clients (all-interfaces | configured-interfaces | non-configured-interfaces);

```

```

group group-name {
  access-profile profile-name;
  active-server-group server-group-name;
  authentication {
    password password-string;
    username-include {
      circuit-type;
      client-id;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-description (device-interface | logical-interface);
      interface-name interface-name;
      logical-system-name;
      mac-address mac-address;
      relay-agent-interface-id;
      relay-agent-remote-id;
      relay-agent-subscriber-id;
      routing-instance-name;
      user-prefix user-prefix-string;
      vlan-tags;
    }
  }
}

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

forward-only {
  logical-system <current | default | logical-system-name>;
  routing-instance <current | default | routing-instance-name>;
}

interface interface-name {
  access-profile profile-name;
  dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
  }
  exclude;
  overrides {
    allow-snooped-clients;
    asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    client-negotiation-match incoming-interface;
    delay-authentication;
    delete-binding-on-renegotiation;
  }
}

```

```

    dual-stack dual-stack-group-name;
    interface-client-limit number;
    no-allow-snooped-clients;
    no-bind-on-request;
    relay-source interface-name;
    send-release-on-delete;
}
service-profile dynamic-profile-name;
short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
trace;
upto upto-interface-name;
}
}
lease-time-validation {
    lease-time-threshold seconds;
    violation-action action;
}
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;
        }
        layer2-liveness-detection {
            max-consecutive-retries number;
            transmit-interval interval;
        }
    }
}
}

```

```

overrides {
    allow-snooped-clients;
    asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    client-negotiation-match incoming-interface;
    delay-authentication;
    delete-binding-on-renegotiation;
    dual-stack dual-stack-group-name;
    interface-client-limit number;
    no-allow-snooped-clients;
    no-bind-on-request;
    relay-source interface-name;
    send-release-on-delete;
}
relay-agent-interface-id {
    include-irb-and-l2;
    keep-incoming-interface-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82 <strict>;
    use-vlan-id;
}
relay-agent-remote-id {
    include-irb-and-l2;
    keep-incoming-interface-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82 <strict>;
    use-vlan-id;
}

```



```

relay-option {
    option-number option-number;
    default-action {
        drop;
        forward-only;
        relay-server-group relay-server-group;
    }
    equals (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        relay-server-group relay-server-group;
    }
    starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        relay-server-group relay-server-group;
    }
}
remote-id-mismatch disconnect;
route-suppression;
service-profile dynamic-profile-name;
short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
leasequery {
    attempts number-of-attempts;
    timeout seconds;
}
lease-time-validation {
    lease-time-threshold seconds;
    violation-action action;
}

```

```

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;
    }
    layer2-liveness-detection {
      max-consecutive-retries number;
      transmit-interval interval;
    }
    route-suppression;
    service-profile dynamic-profile-name;
  }
}

no-snoop;
overrides {
  allow-snooped-clients;
  asymmetric-lease-time seconds;
  asymmetric-prefix-lease-time seconds;
  client-negotiation-match incoming-interface;
  delay-authentication;
  delete-binding-on-renegotiation;
  dual-stack dual-stack-group-name;
  interface-client-limit number;
  no-allow-snooped-clients;
  no-bind-on-request;
  relay-source interface-name;
  send-release-on-delete;
}

```

```

relay-agent-interface-id {
    include-irb-and-l2;
    keep-incoming-interface-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82 <strict>;
    use-vlan-id;
}
relay-agent-remote-id {
    include-irb-and-l2;
    keep-incoming-remote-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82 <strict>;
    use-vlan-id;
}
relay-option {
    option-number option-number;
    default-action {
        drop;
        forward-only;
        relay-server-group relay-server-group;
    }
    equals (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        relay-server-group relay-server-group;
    }
    starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        relay-server-group relay-server-group;
    }
}
relay-option-vendor-specific{
    host-name;
    location;
remote-id-mismatch disconnect;
route-suppression;

```

```
server-group {  
  server-group-name {  
    server-ip-address;  
  }  
}  
server-response-time seconds;  
service-profile dynamic-profile-name;  
short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;  
}
```

```

dual-stack-group dual-stack-group-name {
  access-profile profile-name;
  authentication {
    password password-string;
    username-include {
      circuit-type;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-description (device-interface | logical-interface);
      interface-name;
      logical-system-name;
      mac-address;
      relay-agent-interface-id;
      relay-agent-remote-id;
      routing-instance-name;
      user-prefix user-prefix-string;
      vlan-tags;
    }
  }
  classification-key {
    circuit-id circuit-id;
    mac-address mac-address;
    remote-id remote-id;
  }
  dual-stack-interface-client-limit number;
  dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
  }
  liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
      layer2-liveness-detection {
        max-consecutive-retries number;
        transmit-interval interval;
      }
    }
  }
  protocol-master (inet | inet6);
  relay-agent-interface-id {
    include-irb-and-l2;
    keep-incoming-interface-id ;
    no-vlan-interface-name;
    prefix prefix;
  }
}

```

```

    use-interface-description (logical | device);
    use-option-82 <strict>;
    use-vlan-id;
  }
  relay-agent-remote-id {
    include-irb-and-l2;
    keep-incoming-remote-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82 <strict>;
    use-vlan-id;
  }
  service-profile dynamic-profile-name;
  short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}
duplicate-clients-in-subnet (incoming-interface | option-82):
dynamic-profile profile-name {
  aggregate-clients (merge | replace);
  use-primary primary-profile-name;
}
forward-only {
  logical-system <current | default | logical-system-name>;
  routing-instance <current | default | routing-instance-name>;
}
forward-only-replies;
forward-snooped-clients (all-interfaces | configured-interfaces | non-configured-interfaces);

```

```

group group-name {
  access-profile profile-name;
  active-server-group server-group-name;
  authentication {
    password password-string;
    username-include {
      circuit-type;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-description (device-interface | logical-interface);
      interface-name interface-name;
      logical-system-name;
      mac-address;
      option-60;
      option-82 [circuit-id] [remote-id];
      routing-instance-name;
      user-prefix user-prefix-string;
    }
    vlan-tags;
  }
  dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
  }
  forward-only {
    logical-system <current | default | logical-system-name>;
    routing-instance <current | default | routing-instance-name>;
  }
  forward-only {
    logical-system <current | default | logical-system-name>;
    routing-instance <current | default | routing-instance-name>;
  }
  interface interface-name {
    access-profile profile-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 {
    allow-no-end-option;
    allow-snooped-clients;
    always-write-giaddr;
    always-write-option-82;
    asymmetric-lease-time seconds;
    client-discover-match <option60-and-option82 | incoming-interface>;
    delay-authentication;
    delete-binding-on-renegotiation;
    disable-relay;
    dual-stack dual-stack-group-name;
    interface-client-limit number;
    layer2-unicast-replies;
    no-allow-snooped-clients;
    no-bind-on-request;
    proxy-mode;
    relay-source
    replace-ip-source-with;
    send-release-on-delete;
    trust-option-82;
}

service-profile dynamic-profile-name;
short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
trace;
upto upto-interface-name;
}

```



```

overrides {
    allow-no-end-option
    allow-snooped-clients;
    always-write-giaddr;
    always-write-option-82;
    asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    client-discover-match (option60-and-option82 | incoming-interface);
    delay-authentication;
    delete-binding-on-renegotiation;
    disable-relay;
    dual-stack dual-stack-group-name;
    interface-client-limit number;
    layer2-unicast-replies;
    no-allow-snooped-clients;
    no-bind-on-request;
    proxy-mode;
    relay-source
    replace-ip-source-with;
    send-release-on-delete;
    trust-option-82;
}
relay-option {
    option-number option-number;
    default-action {
        drop;
        forward-only;
        relay-server-group group-name;
    }
    equals (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        relay-server-group relay-server-group;
    }
    starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        local-server-group local-server-group;
        relay-server-group relay-server-group;
    }
}

```

```
relay-option-82 {  
  circuit-id {  
    prefix prefix;  
    use-interface-description (logical | device);  
  }  
  remote-id {  
    prefix prefix;  
    use-interface-description (logical | device);  
  }  
  server-id-override  
}  
remote-id-mismatch disconnect;  
route-suppression:  
service-profile dynamic-profile-name;  
short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;  
}  
leasequery {  
  attempts number-of-attempts;  
  timeout seconds;  
}  
lease-time-validation {  
  lease-time-threshold seconds;  
  violation-action action;  
}
```

```
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;  
    }  
    layer2-liveness-detection {  
      max-consecutive-retries number;  
      transmit-interval interval;  
    }  
  }  
}  
no-snoop;
```

```

overrides {
  allow-no-end-option
  allow-snooped-clients;
  always-write-giaddr;
  always-write-option-82;
  asymmetric-lease-time seconds;
  asymmetric-prefix-lease-time seconds;
  client-discover-match (option60-and-option82 | incoming-interface);
  delay-authentication;
  delete-binding-on-renegotiation;
  disable-relay;
  dual-stack dual-stack-group-name;
  interface-client-limit number;
  layer2-unicast-replies;
  no-allow-snooped-clients;
  no-bind-on-request;
  proxy-mode;
  relay-source
  replace-ip-source-with;
  send-release-on-delete;
  trust-option-82;
}
relay-option {
  option-number option-number;
  default-action {
    drop;
    forward-only;
    relay-server-group group-name;
  }
  equals (ascii ascii-string | hexadecimal hexadecimal-string) {
    drop;
    forward-only;
    relay-server-group relay-server-group;
  }
  starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
    drop;
    forward-only;
    local-server-group local-server-group;
    relay-server-group relay-server-group;
  }
}

```

```

relay-option-82 {
  circuit-id {
    prefix prefix;
    use-interface-description (logical | device);
  }
  remote-id {
    prefix prefix;
    use-interface-description (logical | device);
  }
  server-id-override
}
}
remote-id-mismatch disconnect;
route-suppression:
server-group {
  server-group-name {
    server-ip-address;
  }
}
server-response-time seconds;
service-profile dynamic-profile-name;
short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}

```

Hierarchy Level

```

[edit forwarding-options],
[edit logical-systems logical-system-name forwarding-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options],
[edit routing-instances routing-instance-name forwarding-options]

```

Release Information

Statement introduced in Junos OS Release 8.3.

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

Statement introduced in Junos OS Release 13.2X51 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure extended Dynamic Host Configuration Protocol (DHCP) relay and DHCPv6 relay options on the router or switch to enable the router (or switch) 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.

DHCP relay supports the attachment of dynamic profiles and also interacts with the local AAA Service Framework to use back-end authentication servers, such as RADIUS, to provide subscriber authentication or client authentication. You can attach dynamic profiles and configure authentication support on a global basis or for a specific group of interfaces.

The extended DHCP and DHCPv6 relay agent options configured with the **dhcp-relay** and **dhcpv6** statements are incompatible with the DHCP/BOOTP relay agent options configured with the **bootp** statement. As a result, the extended DHCP or DHCPv6 relay agent and the DHCP/BOOTP relay agent cannot both be enabled on the router (or switch) at the same time.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

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

RELATED DOCUMENTATION

<i>Extended DHCP Relay Agent Overview</i>
<i>DHCPv6 Relay Agent Overview</i>
<i>DHCP Relay Proxy Overview</i>
<i>Specifying Authentication Support</i>

dhcpv6 (DHCP Local Server)

Syntax

```

dhcpv6 {
  access-profile profile-name;
  allow-active-leasequery {
    idle-timeout seconds;
    peer-address address;
    timeout seconds;
  }
  allow-bulk-leasequery {
    max-connections number-of-connections;
    max-empty-replies number-of-replies;
    restricted-requestor;
    timeout seconds;
  }
  allow-leasequery {
    restricted-requestor;
  }
  authentication {
    password password-string;
    username-include {
      circuit-type;
      client-id;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-description (device-interface | logical-interface);
      logical-system-name;
      mac-address;
      relay-agent-interface-id;
      relay-agent-remote-id;
      relay-agent-subscriber-id;
      routing-instance-name;
      user-prefix user-prefix-string;
      vlan-tags;
    }
  }
  duplicate-clients incoming-interface;
  group group-name {
    access-profile profile-name;
    authentication {
      ...
    }
    interface interface-name {
      access-profile profile-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 {
  asymmetric-lease-time seconds;
  asymmetric-prefix-lease-time seconds;
  client-negotiation-match incoming-interface;
  delay-advertise {
    based-on (option-15 | option-16 | option-18 | option-37) {
      equals {
        ascii ascii-string;
        hexadecimal hexadecimal-string;
      }
      not-equals {
        ascii ascii-string;
        hexadecimal hexadecimal-string;
      }
      starts-with {
        ascii ascii-string;
        hexadecimal hexadecimal-string;
      }
    }
    delay-time seconds;
  }
  delete-binding-on-renegotiation;
  interface-client-limit number;
  multi-address-embedded-option-response;
  process-inform {
    pool pool-name;
  }
  protocol-attributes attribute-set-name;
  rapid-commit;
}
service-profile dynamic-profile-name;
short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
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;  
    }  
    layer2-liveness-detection {  
      max-consecutive-retries number;  
      transmit-interval interval;  
    }  
  }  
}
```

```

overrides {
  asymmetric-lease-time seconds;
  asymmetric-prefix-lease-time seconds;
  client-negotiation-match incoming-interface;
  delay-advertise {
    based-on (option-15 | option-16 | option-18 | option-37) {
      equals {
        ascii ascii-string;
        hexadecimal hexadecimal-string;
      }
      not-equals {
        ascii ascii-string;
        hexadecimal hexadecimal-string;
      }
      starts-with {
        ascii ascii-string;
        hexadecimal hexadecimal-string;
      }
    }
    delay-time seconds;
  }
  delegated-pool;
  delete-binding-on-renegotiation;
  interface-client-limit number;
  multi-address-embedded-option-response;
  process-inform {
    pool pool-name;
  }
  protocol-attributes attribute-set-name;
  rapid-commit;
}
route-suppression;
service-profile dynamic-profile-name;
short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}

```

```
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;  
    }  
    layer2-liveness-detection {  
      max-consecutive-retries number;  
      transmit-interval interval;  
    }  
  }  
}
```

```

overrides {
  asymmetric-lease-time seconds;
  asymmetric-prefix-lease-time seconds;
  client-negotiation-match incoming-interface;
  delay-advertise {
    based-on (option-15 | option-16 | option-18 | option-37) {
      equals {
        ascii ascii-string;
        hexadecimal hexadecimal-string;
      }
      not-equals {
        ascii ascii-string;
        hexadecimal hexadecimal-string;
      }
      starts-with {
        ascii ascii-string;
        hexadecimal hexadecimal-string;
      }
    }
    delay-time seconds;
  }
  delegated-pool;
  delete-binding-on-renegotiation;
  interface-client-limit number;
  multi-address-embedded-option-response;
  process-inform {
    pool pool-name;
  }
  protocol-attributes attribute-set-name;
  rapid-commit;
  reconfigure {
    attempts attempt-count;
    clear-on-abort;
    strict;
    timeout timeout-value;
    token token-value;
    trigger {
      radius-disconnect;
    }
  }
}
reauthenticate (<lease-renewal> <remote-id-mismatch >);

```

```

reconfigure {
  attempts attempt-count;
  clear-on-abort;
  strict;
  support-option-pd-exclude;
  timeout timeout-value;
  token token-value;
  trigger {
    radius-disconnect;
  }
}
requested-ip-network-match subnet-mask;
route-suppression;
server-duid-type type;
service-profile dynamic-profile-name;
short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}

```

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 9.6.

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

Description

Configure DHCPv6 local server options on the router or switch to enable the router or switch to function as a server for the DHCP protocol for IPv6. The DHCPv6 local server sends and receives packets using the IPv6 protocol and informs IPv6 of the routing requirements of router clients. The local server works together with the AAA service framework to control subscriber access (or DHCP client access) and accounting.

The DHCPv6 local server is fully compatible with the extended DHCP local server and DHCP relay agent.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| *DHCPv6 Local Server Overview*

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 dual-stack-group dual-stack-group-name],
[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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces

Configuring a Default Subscriber Service

dynamic-profile (DHCP Relay Agent)

Syntax

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

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface interface-name],
[edit forwarding-options dhcp-relay dual-stack-group dual-stack-group-name],
[edit forwarding-options dhcp-relay group group-name],
[edit forwarding-options dhcp-relay group group-name interface interface-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Release Information

Statement introduced in Junos OS Release 9.2.

Support at the **[edit ... dhcpv6]** hierarchy levels introduced in Junos OS Release 11.4.

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

Support at the **[edit ... dual-stack-group *dual-stack-group-name*]** hierarchy level introduced in Junos OS Release 15.1.

Description

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

M120 and M320 routers do not support DHCPv6.

Options

profile-name—Name of the dynamic profile.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[dhcp-relay](#) | 264

Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces

Grouping Interfaces with Common DHCP Configurations

Configuring a Default Subscriber Service

dynamic-profile (Dynamic PPPoE)

Syntax

```
dynamic-profile profile-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family pppoe],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family pppoe],
[edit interfaces interface-name unit logical-unit-number family pppoe],
[edit interfaces interface-name unit logical-unit-number pppoe-underlying-options],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family pppoe],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number pppoe-underlying-options]
```

Release Information

Statement introduced in Junos OS Release 10.1.

Support for the **[edit ... family pppoe]** hierarchies introduced in Junos OS Release 11.2.

Description

Attach a PPPoE dynamic profile to an underlying Ethernet interface. This underlying interface is configured with either the **encapsulation ppp-over-ether** statement or the **family pppoe** statement; the two statements are mutually exclusive. When the router creates a dynamic PPPoE logical interface on the underlying interface, it uses the information in the dynamic profile to determine the properties of the dynamic PPPoE logical interface.

NOTE: The **[edit ... family pppoe]** hierarchies are supported only on MX Series routers with MPCs.

Starting in Junos OS Release 17.2R1, you can configure converged services for MS-MPCs and MS-MICs. You can configure captive portal content delivery (CPCD) profiles for MS-MICs and MS-MPCs by including the service interface ms-fpc/pic/port statement at the **edit service-set service set name captive-portal-content-delivery-profile *profile name* interface-service** heirarchy level.

Options

profile-name—Name of a previously configured PPPoE dynamic profile, up to 64 characters in length, defined at the **[edit dynamic-profiles *profile-name* interfaces pp0]** hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces

Configuring the PPPoE Family for an Underlying Interface

Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview

dynamic-profile (Stacked VLAN)

Syntax

```
dynamic-profile profile-name {
  accept (any | dhcp-v4 | dhcp-v6 | inet | inet6 | pppoe);
  access-profile vlan-dynamic-profile-name;
  ranges (any | low-tag-high-tag), (any | low-tag-high-tag);
}
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure stacked-vlan-ranges]
```

Release Information

Statement introduced in Junos OS Release 9.5.

Description

Configure a dynamic profile for use when configuring dynamic stacked VLANs.

Options

profile-name—Name of the dynamic profile that you want to use when configuring dynamic stacked VLANs.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Dynamic Profiles Overview](#)

[Configuring a Basic Dynamic Profile](#)

[Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs](#)

dynamic-profile (VLAN)

Syntax

```
dynamic-profile profile-name {
  accept (any | dhcp-v4 | dhcp-v6 | inet | inet6 | pppoe);
  accept-out-of-band protocol;
  access-profile vlan-dynamic-profile-name;
  ranges (any | low-tag)-(any | high-tag);
}
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges]
```

Release Information

Statement introduced in Junos OS Release 9.5.

Description

Configure a dynamic profile for use when configuring dynamic VLANs.

Options

profile-name—Name of the dynamic profile that you want to use when configuring dynamic VLANs.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Dynamic Profiles Overview

Configuring a Basic Dynamic Profile

Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs

dynamic-profiles

Syntax

```
dynamic-profiles {
  profile-name {
    class-of-service {
      dynamic-class-of-service-options {
        vendor-specific-tags tag;
      }
    }
    interfaces {
      interface-name ;
    }
    unit logical-unit-number {
      classifiers {
        type (classifier-name | default);
      }
      output-traffic-control-profile (profile-name | $junos-cos-traffic-control-profile);
      report-ingress-shaping-rate bps;
      rewrite-rules {
        dscp (rewrite-name | default);
        dscp-ipv6 (rewrite-name | default);
        ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
        inet-precedence (rewrite-name | default);
      }
    }
  }
}

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

schedulers {
  (scheduler-name) {
    buffer-size (seconds | percent percentage | remainder | temporal microseconds);
    drop-profile-map loss-priority (any | low | medium-low | medium-high | high) protocol (any | non-tcp | tcp) drop-profile profile-name;
    excess-priority (low | high | $junos-cos-scheduler-excess-priority);
    excess-rate (percent percentage | percent $junos-cos-scheduler-excess-rate);
    overhead-accounting (shaping-mode) <bytes (byte-value)>;
    priority priority-level;
    shaping-rate (rate | predefined-variable);
    transmit-rate (percent percentage | rate | remainder) <exact | rate-limit>;
  }
}
```



```

}
traffic-control-profiles profile-name {
    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>;
    max-burst-size cells;
    overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
    peak-rate rate;
    scheduler-map map-name;
    shaping-rate (percent percentage | rate | predefined-variable) <burst-size bytes>;
    shaping-rate-excess-high (percent percentage | rate) <burst-size bytes>;
    shaping-rate-excess-medium-high (percent percentage | rate) <burst-size bytes>;
    shaping-rate-excess-medium-low (percent percentage | rate) <burst-size bytes>;
    shaping-rate-excess-low (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-high (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-low (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-medium (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-medium-low (percent percentage | rate) <burst-size bytes>;
    shaping-rate-priority-strict-high (percent percentage | rate) <burst-size bytes>;
    sustained-rate rate;
}
}

```

```

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 {
      enhanced-mode-override;
      fast-lookup-filter;
      instance-shared;
      interface-shared;
    }
    interface-specific;
    term term-name {
      from {
        match-conditions;
      }
      then {
        action;
        action-modifiers;
      }
      only-at-create;
    }
    filter filter-name {
      interface-specific;
      term term-name {
        from {
          match-conditions;
        }
        then {
          action;
          action-modifiers;
        }
      }
    }
  }
  hierarchical-policer uid {
    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;
    }
}
}
policer uid {
    filter-specific;
    if-exceeding {
        (bandwidth-limit bps | bandwidth-percent percentage);
        burst-size-limit bytes;
    }
    logical-bandwidth-policer;
    logical-interface-policer;
    physical-interface-policer;
    then {
        policer-action;
    }
}
}

```

```
three-color-policer uid {  
    action {  
        loss-priority high then discard;  
    }  
    logical-interface-policer;  
    single-rate {  
        (color-aware | color-blind);  
        committed-burst-size bytes;  
        committed-information-rate bps;  
        excess-burst-size bytes;  
    }  
    two-rate {  
        (color-aware | color-blind);  
        committed-burst-size bytes;  
        committed-information-rate bps;  
        peak-burst-size bytes;  
        peak-information-rate bps;  
    }  
}  
}
```

```

interfaces interface-name {
  interface-set interface-set-name {
    interface interface-name {
      unit logical unit number {
        advisory-options {
          downstream-rate rate;
          upstream-rate rate;
        }
      }
    }
  }
}

unit logical-unit-number {
  actual-transit-statistics;
  auto-configure {
    agent-circuit-identifier {
      dynamic-profile profile-name;
    }
    line-identity {
      include {
        accept-no-ids;
        circuit-id;
        remote-id;
      }
      dynamic-profile profile-name;
    }
  }
}

encapsulation (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-cisco-nlpid | atm-tcc-vc-mux | atm-mlppp-llc |
atm-nlpid | atm-ppp-llc | atm-ppp-vc-mux | atm-snap | atm-tcc-snap | atm-vc-mux | ether-over-atm-llc
| ether-vpls-over-atm-llc | ether-vpls-over-fr | ether-vpls-over-ppp | ethernet | frame-relay-ccc |
frame-relay-ppp | frame-relay-tcc | frame-relay-ether-type | frame-relay-ether-type-tcc |
multilink-frame-relay-end-to-end | multilink-ppp | ppp-over-ether | ppp-over-ether-over-atm-llc |
vlan-bridge | vlan-ccc | vlan-vci-ccc | vlan-tcc | vlan-vpls);

family family {
  address address;
  filter {
    adf {
      counter;
      input-precedence precedence;
      not-mandatory;
      output-precedence precedence;
      rule rule-value;
    }
    input filter-name (
      precedence precedence;

```

```

        shared-name filter-shared-name;
    }
    output filter-name {
        precedence precedence;
        shared-name filter-shared-name;
    }
}
rpf-check {
    fail-filter filter-name;
    mode loose;
}
service {
    input {
        service-set service-set-name {
            service-filter filter-name;
        }
        post-service-filter filter-name;
    }
    input-vlan-map {
        inner-tag-protocol-id tpid;
        inner-vlan-id number;
        (push | swap);
        tag-protocol-id tpid;
        vlan-id number;
    }
    output {
        service-set service-set-name {
            service-filter filter-name;
        }
    }
    output-vlan-map {
        inner-tag-protocol-id tpid;
        inner-vlan-id number;
        (pop | swap);
        tag-protocol-id tpid;
        vlan-id number;
    }
    pcef pcef-profile-name {
        activate rule-name | activate-all;
    }
}
unnumbered-address interface-name <preferred-source-address address>;
}

```

```

filter {
    input filter-name (
        shared-name filter-shared-name;
    )
    output filter-name {
        shared-name filter-shared-name;
    }
}
host-prefix-only;
ppp-options {
    aaa-options aaa-options-name;
    authentication [ authentication-protocols ];
    chap {
        challenge-length minimum minimum-length maximum maximum-length;
        local-name name;
    }
    ignore-magic-number-mismatch;
    initiate-ncp (dual-stack-passive | ipv6 | ip)
    ipcp-suggest-dns-option;
    mru size;
    mtu (size | use-lower-layer);
    on-demand-ip-address;
    pap;
    peer-ip-address-optional;
    local-authentication {
        password password;
        username-include {
            circuit-id;
            delimiter character;
            domain-name name;
            mac-address;
            remote-id;
        }
    }
}
targeted-options {
    backup backup;
    group group;
    primary primary;
    weight ($junos-interface-target-weight | weight-value);
}

```

```

telemetry {
  subscriber-statistics;
  queue-statistics {
    interface $junos-interface-name {
      refresh rate;
      queues queue set;
    }
    interface-set $junos-interface-set-name {
      refresh rate;
      queues queue set;
    }
  }
}
vlan-id number;
vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
}
}
interfaces {
  demux0 {...}
}
interfaces {
  pp0 {...}
}
policy-options {
  prefix-list uid {
    ip-addresses;
    dynamic-db;
  }
}
predefined-variable-defaults predefined-variable <variable-option> default-value;
profile-type remote-device-service;

```



```

protocols {
  igmp {
    interface interface-name {
      accounting;
      disable;
      group-limit limit;
      group-policy;
      group-threshold value;
      immediate-leave
      log-interval seconds;
      no-accounting;
      oif-map;
      passive;
      promiscuous-mode;
      ssm-map ssm-map-name;
      ssm-map-policy ssm-map-policy-name
      static {
        group group {
          source source;
        }
      }
      version version;
    }
  }
  mld {
    interface interface-name {
      (accounting | no-accounting);
      disable;
      group-limit limit;
      group-policy;
      group-threshold value;
      immediate-leave;
      log-interval seconds;
      oif-map;
      passive;
      ssm-map ssm-map-name;
      ssm-map-policy ssm-map-policy-name;
      static {
        group multicast-group-address {
          exclude;
          group-count number;
          group-increment increment;
          source ip-address {
            source-count number;

```

```
        source-increment increment;  
    }  
}  
}  
version version;  
}  
}  
router-advertisement {  
    interface interface-name {  
        current-hop-limit number;  
        default-lifetime seconds;  
        (managed-configuration | no-managed-configuration);  
        max-advertisement-interval seconds;  
        min-advertisement-interval seconds;  
        (other-stateful-configuration | no-other-stateful-configuration);  
        prefix prefix;  
        reachable-time milliseconds;  
        retransmit-timer milliseconds;  
    }  
}  
}
```

```

routing-instances routing-instance-name {
  interface interface-name;
  routing-options {
    access {
      route prefix {
        next-hop next-hop;
        metric route-cost;
        preference route-distance;
        tag route-tag;
        tag2 route-tag2;
      }
    }
    access-internal {
      route subscriber-ip-address {
        qualified-next-hop underlying-interface {
          mac-address address;
        }
      }
    }
    multicast {
      interface interface-name {
        no-qos-adjust;
      }
    }
  }
  rib routing-table-name {
    access {
      route prefix {
        next-hop next-hop;
        metric route-cost;
        preference route-distance;
        tag route-tag;
        tag2 route-tag2;
      }
    }
    access-internal {
      route subscriber-ip-address {
        qualified-next-hop underlying-interface {
          mac-address address;
        }
      }
    }
  }
}

```

```

routing-options {
  access {
    route prefix {
      next-hop next-hop;
      metric route-cost;
      preference route-distance;
      tag route-tag;
      tag2 route-tag2;
    }
  }
  access-internal {
    route subscriber-ip-address {
      qualified-next-hop underlying-interface {
        mac-address address;
      }
    }
  }
  multicast {
    interface interface-name {
      no-qos-adjust;
    }
  }
}

services {
  captive-portal-content-delivery {
    auto-deactivate value;
    rule name {
      match-direction (input | input-output | output);
      term name {
        then {
          accept;
          redirect url;
          rewrite destination-address address <destination-port port-number>;
          syslog;
        }
      }
    }
  }
}

```

```

variables {
  variable-name {
    default-value default-value;
    equals expression;
    mandatory;
    uid;
    uid-reference;
  }
}

```

Hierarchy Level

[edit]

Release Information

Statement introduced in Junos OS Release 9.2.

Support at the **filter**, **policer**, **hierarchical-policer**, **three-color-policer**, and **policy options** hierarchy levels introduced in Junos OS Release 11.4.

Description

Create dynamic profiles for use with DHCP or PPP client access.

Options

profile-name—Name of the dynamic profile; string of up to 80 alphanumeric characters.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring a Basic Dynamic Profile

Configuring Dynamic VLANs Based on Agent Circuit Identifier Information

Dynamic Profiles for Subscriber Management

egress-stats (Flat-File Accounting Options)

Syntax

```
egress-stats {  
    all-fields;  
    input-bytes;  
    input-packets;  
    output-bytes;  
    output-packets;  
    queue-id;  
    red-drop-bytes;  
    red-drop-packets;  
    tail-drop-packets;  
}
```

Hierarchy Level

```
[edit accounting-options flat-file-profile profile-name fields]
```

Release Information

Statement introduced in Junos OS Release 16.1R4.

Description

Specify egress queue statistics to be collected for the interface.

Options

all-fields—Collect all egress queue statistics available for the interface context, logical or physical.

input-bytes—Collect the number of octets queued including traffic dropped because of congestion.

input-packets—Collect the number of packets queued including traffic dropped because of congestion.

output-bytes—Collect the number of octets transmitted by the egress queue.

output-packets—Collect the number of packets transmitted by the egress queue.

queue-id—Collect the logical identifier for the egress queue; identifies the traffic class.

red-drop-bytes—Collect the number of octets dropped on the egress queue because of random early detection.

red-drop-packets—Collect the number of packets dropped on the egress queue because of random early detection.

tail-drop-packets—Collect the number of packets dropped in the egress queue because of tail drop.

Required Privilege Level

system—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Flat-File Accounting for Layer 2 Wholesale | 153](#)

[Configuring Flat-File Accounting for Extensible Subscriber Services Management | 158](#)

[Flat-File Accounting Overview | 149](#)

encapsulation (Dynamic Interfaces)

Syntax

```
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};
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 10.4.

Description

Dynamic interface configuration of the logical link-layer encapsulation type.

Options

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

atm-ccc-vc-mux—Use ATM virtual circuit (VC) multiplex encapsulation on circuit cross-connect (CCC) circuits. When you use this encapsulation type, you can configure the **ccc** family only.

atm-cisco-nlpid—Use Cisco ATM network layer protocol ID (NLPID) encapsulation. When you use this encapsulation type, you can configure the **inet** family only.

atm-mlppp-llc—For ATM2 IQ interfaces only, use Multilink Point-to-Point Protocol (MLPPP) over AAL5 LLC. For this encapsulation type, your router must be equipped with a link services or voice services PIC. MLPPP over ATM encapsulation is not supported on ATM2 IQ OC48 interfaces.

atm-nlpid—Use ATM NLPID encapsulation. When you use this encapsulation type, you can configure the **inet** family only.

atm-ppp-llc—For ATM2 IQ interfaces only, use PPP over AAL5 LLC encapsulation.

atm-ppp-vc-mux—For ATM2 IQ interfaces only, use PPP over ATM AAL5 multiplex encapsulation.

atm-snap—Use ATM subnetwork attachment point (SNAP) encapsulation.

atm-tcc-snap—Use ATM SNAP encapsulation on translational cross-connect (TCC) circuits.

atm-tcc-vc-mux—Use ATM VC multiplex encapsulation on TCC circuits. When you use this encapsulation type, you can configure the **tcc** family only.

atm-vc-mux—Use ATM VC multiplex encapsulation. When you use this encapsulation type, you can configure the **inet** family only.

ether-over-atm-llc—For interfaces that carry IPv4 traffic, use Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces.

ether-vpls-over-atm-llc—For ATM2 IQ interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.

ether-vpls-over-fr—For E1, T1, E3, T3, and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over Frame Relay encapsulation to support Bridged Ethernet over Frame Relay encapsulated TDM interfaces for VPLS applications, as per *Multiprotocol Interconnect over Frame Relay* (RFC 2427 [1490]).

ether-vpls-over-ppp—For E1, T1, E3, T3 and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over PPP encapsulation to support Bridged Ethernet over PPP encapsulated TDM interfaces for VPLS applications.

ethernet—Use Ethernet II encapsulation (as described in RFC 894, *A Standard for the Transmission of IP Datagrams over Ethernet Networks*).

ethernet-vpls—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard Tag Protocol ID (TPID) values.

extended-vlan-vpls—Use extended virtual LAN (VLAN) VPLS encapsulation on Ethernet interfaces that have VLAN 802.1Q tagging and VPLS enabled and that must accept packets carrying TPIDs 0x8100, 0x9100, and 0x9901.

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

frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

frame-relay-ppp—Use PPP over Frame Relay circuits. When you use this encapsulation type, you can configure the **ppp** family only.

frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits for connecting unlike media. When you use this encapsulation type, you can configure the **tcc** family only.

frame-relay-ether-type—Use Frame Relay ether type encapsulation for compatibility with Cisco Frame Relay. The physical interface must be configured with **flexible-frame-relay** encapsulation.

frame-relay-ether-type-tcc—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect unlike media. The physical interface must be configured with **flexible-frame-relay** encapsulation.

multilink-frame-relay-end-to-end—Use MLFR FRF.15 encapsulation. This encapsulation is used only on multilink, link services, and voice services interfaces and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.

multilink-ppp—Use MLPPP encapsulation. This encapsulation is used only on multilink, link services, and voice services interfaces and their constituent T1 or E1 interfaces.

ppp-over-ether—You use PPP over Ethernet encapsulation to configure an underlying Ethernet interface for a dynamic PPPoE logical interface.

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

vlan-ccc—Use Ethernet virtual LAN (VLAN) encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

vlan-vci-ccc—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

vlan-tcc—Use Ethernet VLAN encapsulation on TCC circuits. When you use this encapsulation type, you can configure the **tcc** family only.

vlan-vpls—Use Ethernet VLAN encapsulation on VPLS circuits.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring a Retail Dynamic Profile for Use in the Layer 2 Wholesale Solution | 92](#)

[Configuring PPP over ATM2 Encapsulation Overview](#)

exclude (RADIUS Attributes)

Syntax

```
exclude {
  acc-aggr-cir-id-asc [ access-request | accounting-start | accounting-stop ];
  acc-aggr-cir-id-bin [ access-request | accounting-start | accounting-stop ];
  acc-loop-cir-id [ access-request | accounting-start | accounting-stop ];
  acc-loop-encap [ access-request | accounting-start | accounting-stop ];
  acc-loop-remote-id [ access-request | accounting-start | accounting-stop ];
  accounting-authentic [ accounting-off | accounting-on | accounting-start | accounting-stop ];
  accounting-delay-time [ accounting-off | accounting-on | accounting-start | accounting-stop ];
  accounting-session-id access-request;
  accounting-terminate-cause accounting-off;
  acct-request-reason [ accounting-start | accounting-stop ];
  acct-tunnel-connection [ access-request | accounting-start | accounting-stop ];
  act-data-rate-dn [ access-request | accounting-start | accounting-stop ];
  act-data-rate-up [ access-request | accounting-start | accounting-stop ];
  act-interlv-delay-dn [ access-request | accounting-start | accounting-stop ];
  act-interlv-delay-up [ access-request | accounting-start | accounting-stop ];
  att-data-rate-dn [ access-request | accounting-start | accounting-stop ];
  att-data-rate-up [ access-request | accounting-start | accounting-stop ];
  called-station-id [ access-request | accounting-start | accounting-stop ];
  calling-station-id [ access-request | accounting-start | accounting-stop ];
  chargeable-user-identity access-request;
  class [ accounting-start | accounting-stop ];
  cos-shaping-rate [ accounting-start | accounting-stop ];
  delegated-ipv6-prefix [ accounting-start | accounting-stop ];
  dhcp-gi-address [ access-request | accounting-start | accounting-stop ];
  dhcp-header access-request;
  dhcp-mac-address [ access-request | accounting-start | accounting-stop ];
  dhcp-options [ access-request | accounting-start | accounting-stop ];
  dhcpv6-header access-request;
  dhcpv6-options [ access-request | accounting-start | accounting-stop ];
  downstream-calculated-qos-rate [ access-request | accounting-start | accounting-stop ];
  dsl-forum-attributes [ access-request | accounting-start | accounting-stop ];
  dsl-line-state [ access-request | accounting-start | accounting-stop ];
  dsl-type [ access-request | accounting-start | accounting-stop ];
  dynamic-iflset-name [ accounting-start | accounting-stop ];
  event-timestamp [ accounting-off | accounting-on | accounting-start | accounting-stop ];
  filter-id [ accounting-start | accounting-stop ];
  first-relay-ipv4-address [ access-request | accounting-start | accounting-stop ];
  first-relay-ipv6-address [ access-request | accounting-start | accounting-stop ];
  framed-interface-id [ access-request | accounting-start | accounting-stop ];
  framed-ip-address [ access-request | accounting-start | accounting-stop ];
```

```

framed-ip-netmask [ access-request | accounting-start | accounting-stop ];
framed-ip-route [ accounting-start | accounting-stop ];
framed-ipv6-address [ access-request | accounting-start | accounting-stop ];
framed-ipv6-pool [ accounting-start | accounting-stop ];
framed-ipv6-prefix [ accounting-start | accounting-stop ];
framed-ipv6-route [ accounting-start | accounting-stop ];
framed-pool [ accounting-start | accounting-stop ]; input-ipv6-gigawords accounting-stop;
input-filter [ accounting-start | accounting-stop ];
input-gigapackets accounting-stop;
input-gigawords accounting-stop;
input-ipv6-octets accounting-stop;
input-ipv6-packets accounting-stop;
interface-description [ access-request | accounting-start | accounting-stop ];
l2c-downstream-data [ access-request | accounting-start | accounting-stop ];
l2c-upstream-data [ access-request | accounting-start | accounting-stop ];
l2tp-rx-connect-speed [ access-request | accounting-start | accounting-stop ];
l2tp-tx-connect-speed [ access-request | accounting-start | accounting-stop ];
max-data-rate-dn [ access-request | accounting-start | accounting-stop ];
max-data-rate-up [ access-request | accounting-start | accounting-stop ];
max-interlv-delay-dn [ access-request | accounting-start | accounting-stop ];
max-interlv-delay-up [ access-request | accounting-start | accounting-stop ];
min-data-rate-dn [ access-request | accounting-start | accounting-stop ];
min-data-rate-up [ access-request | accounting-start | accounting-stop ];
min-lp-data-rate-dn [ access-request | accounting-start | accounting-stop ];
min-lp-data-rate-up [ access-request | accounting-start | accounting-stop ];
nas-identifier [ access-request | accounting-off | accounting-on | accounting-start | accounting-stop ];
nas-port [ access-request | accounting-start | accounting-stop ];
nas-port-id [ access-request | accounting-start | accounting-stop ];
nas-port-type [ access-request | accounting-start | accounting-stop ];
output-filter [ accounting-start | accounting-stop ];
output-gigapackets accounting-stop;
output-gigawords accounting-stop;
output-ipv6-gigawords accounting-stop;
output-ipv6-octets accounting-stop;
output-ipv6-packets accounting-stop;
pppoe-description [ access-request | accounting-start | accounting-stop ];
standard-attribute number {
    packet-type [ access-request | accounting-off | accounting-on | accounting-start | accounting-stop ];
}
tunnel-assignment-id [ access-request | accounting-start | accounting-stop ];
tunnel-client-auth-id [ access-request | accounting-start | accounting-stop ];
tunnel-client-endpoint [ access-request | accounting-start | accounting-stop ];
tunnel-medium-type [ access-request | accounting-start | accounting-stop ];
tunnel-server-auth-id [ access-request | accounting-start | accounting-stop ];

```

```

tunnel-server-endpoint [ access-request | accounting-start | accounting-stop ];
tunnel-type [ access-request | accounting-start | accounting-stop ];
upstream-calculated-qos-rate [ access-request | accounting-start | accounting-stop ];
vendor-id id-number {
    vendor-attribute vsa-number {
        packet-type [ access-request | accounting-off | accounting-on | accounting-start | accounting-stop ];
    }
}
virtual-router [ access-request | accounting-start | accounting-stop ];
}

```

Hierarchy Level

[edit access profile *profile-name* radius attributes]

Release Information

Statement introduced in Junos OS Release 9.1.

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

downstream-calculated-qos-rate, **dsl-forum-attributes**, and **upstream-calculated-qos-rate** options added in Junos OS Release 11.4.

cos-shaping-rate and **filter-id** options added in Junos OS Release 13.2.

pppoe-description option added in Junos OS Release 14.2.

virtual-router option added in Junos OS Release 15.1.

first-relay-ipv4-address and **first-relay-ipv6-address** options added in Junos OS Release 16.1.

acc-loop-encap and **acc-loop-remote-id** options added in Junos OS Release 16.1R4.

access-request option support for all tunnel attributes added in Junos OS Release 15.1R7, 16.1R5, 16.2R2, 17.1R2, 17.2R2, and 17.3R1 for MX Series.

packet-type, **standard-attribute**, **vendor-attribute**, and **vendor-id** options added in Junos OS Release 18.1R1.

Description

Configure the router or switch to exclude the specified attributes from being sent in the specified type of RADIUS message. Exclusion can be useful, for example, for attributes that do not change values over the lifetime of a subscriber. By not sending these attributes, you reduce the packet size without losing information. Contrast this behavior with that provided by the **ignore** statement.

You can specify attribute exclusion for multiple RADIUS message types by enclosing the message types, separated by spaces, within brackets ([]). You do not need brackets when specifying a single message type.

Starting in Junos OS Release 18.1R1, you can specify standard RADIUS attributes with the attribute number. You can specify VSAs with the IANA-assigned vendor ID and the VSA number. With this flexible configuration method, you can configure any standard attribute and VSA supported by your platform to be excluded. The configuration has no effect if you configure unsupported attributes, vendors, and VSAs.

The legacy method allows you to configure only those attributes and VSAs for which the statement syntax includes a specific option. Consequently, you can use the legacy method to exclude only a subset of all attributes that can be received in Access-Accept messages.

Not all attributes are available in all types of RADIUS messages.

NOTE: If you exclude an attribute from Acct-Off messages, the attributes are then excluded from Interim-Acct messages.

NOTE: VSAs with dedicated option names include Juniper Networks (IANA vendor ID 4874) and DSL Forum (vendor ID 3561) VSAs.

Options

RADIUS attribute—RADIUS standard attribute or VSA:

- **acc-aggr-cir-id-asc**—Exclude Juniper Networks VSA 26-112, Acc-Aggr-Cir-Id-Asc.
- **acc-aggr-cir-id-bin**—Exclude Juniper Networks VSA 26-111, Acc-Aggr-Cir-Id-Bin.
- **acc-loop-cir-id**—Exclude Juniper Networks VSA 26-110, Acc-Loop-Cir-Id.
- **acc-loop-encap**—Exclude Juniper Networks VSA 26-183, Acc-Loop-Encap.
- **acc-loop-remote-id**—Exclude Juniper Networks VSA 26-182, Acc-Loop-Remote-Id.
- **accounting-authentic**—Exclude RADIUS attribute 45, Acct-Authentic.
- **accounting-delay-time**—Exclude RADIUS attribute 41, Acct-Delay-Time.
- **accounting-session-id**—Exclude RADIUS attribute 44, Acct-Session-Id.
- **accounting-terminate-cause**—Exclude RADIUS attribute 49, Acct-Terminate-Cause.
- **acct-request-reason**—Exclude Juniper Networks VSA 26-210, Acct-Request-Reason.
- **acct-tunnel-connection**—Exclude RADIUS attribute 68, Acct-Tunnel-Connection.
- **act-data-rate-dn**—Exclude Juniper Networks VSA 26-114, Act-Data-Rate-Dn.
- **act-data-rate-up**—Exclude Juniper Networks VSA 26-113, Act-Data-Rate-Up.
- **act-interlv-delay-dn**—Exclude Juniper Networks VSA 26-126, Act-Interlv-Delay-Dn.
- **act-interlv-delay-up**—Exclude Juniper Networks VSA 26-124, Act-Interlv-Delay-Up.
- **att-data-rate-dn**—Exclude Juniper Networks VSA 26-118, Att-Data-Rate-Dn.
- **att-data-rate-up**—Exclude Juniper Networks VSA 26-117, Att-Data-Rate-Up.
- **called-station-id**—Exclude RADIUS attribute 30, Called-Station-Id.
- **calling-station-id**—Exclude RADIUS attribute 31, Calling-Station-Id.
- **chargeable-user-identity**—Exclude RADIUS attribute 89, Chargeable-User-Identity.
- **class**—Exclude RADIUS attribute 25, Class.
- **cos-shaping-rate**—Exclude Juniper Networks VSA 26-177, Cos-Shaping-Rate.
- **delegated-ipv6-prefix**—Exclude RADIUS attribute 123, Delegated-IPv6-Prefix.
- **dhcp-gi-address**—Exclude Juniper Networks VSA 26-57, DHCP-GI-Address.
- **dhcp-header**—Exclude Juniper Networks VSA 26-208, DHCP-Header.
- **dhcp-mac-address**—Exclude Juniper Networks VSA 26-56, DHCP-MAC-Address.
- **dhcp-options**—Exclude Juniper Networks VSA 26-55, DHCP-Options.
- **dhcpv6-header**—Exclude Juniper Networks VSA 26-209, DHCPv6-Header.
- **dhcpv6-options**—Exclude Juniper Networks VSA 26-207, DHCPv6-Options.

- **dynamic-iflset-name**—Exclude Juniper Networks VSA 26-130, Qos-Set-Name.
- **downstream-calculated-qos-rate**—Exclude Juniper Networks VSA 26-141.
- **dsl-forum-attributes**—Exclude DSL Forum VSA (vendor ID 3561) as described in RFC 4679, *DSL Forum Vendor-Specific RADIUS Attributes*.
- **dsl-line-state**—Exclude Juniper Networks VSA 26-127, DSL-Line-State.
- **dsl-type**—Exclude Juniper Networks VSA 26-128, DSL-Type.
- **event-timestamp**—Exclude RADIUS attribute 55, Event-Timestamp.
- **filter-id**—Exclude RADIUS attribute 11, Filter-Id.
- **first-relay-ipv4-address** —Exclude Juniper Networks VSA 26-189, DHCP-First-Relay-IPv4-Address.
- **first-relay-ipv6-address** —Exclude Juniper Networks VSA 26-190, DHCP-First-Relay-IPv6-Address.
- **framed-interface-id**—Exclude RADIUS attribute 96, Framed-Interface-ID.
- **framed-ip-address**—Exclude RADIUS attribute 8, Framed-IP-Address.
- **framed-ip-netmask**—Exclude RADIUS attribute 9, Framed-IP-Netmask.
- **framed-ip-route**—Exclude RADIUS attribute 22, Framed-Route.
- **framed-ipv6-address**—Exclude RADIUS attribute 168, Framed-IPv6-Address.
- **framed-ipv6-pool**—Exclude RADIUS attribute 100, Framed-IPv6-Pool.
- **framed-ipv6-prefix**—Exclude RADIUS attribute 97, Framed-IPv6-Prefix.
- **framed-ipv6-route**—Exclude RADIUS attribute 99, Framed-IPv6-Route.
- **framed-pool**—Exclude RADIUS attribute 88, Framed-Pool.
- **input-filter**—Exclude Juniper Networks VSA 26-10, Ingress-Policy-Name.
- **input-gigapackets**—Exclude Juniper Networks VSA 26-42, Acct-Input-Gigapackets.
- **input-gigawords**—Exclude RADIUS attribute 52, Acct-Input-Gigawords.
- **input-ipv6-gigawords**—Exclude Juniper Networks VSA 26-155, Acct-Input-IPv6-Gigawords.
- **input-ipv6-octets**—Exclude Juniper Networks VSA 26-151, Acct-Input-IPv6-Octets.
- **input-ipv6-packets**—Exclude Juniper Networks VSA 26-153, Acct-Input-IPv6-Packets.
- **interface-description**—Exclude Juniper Networks VSA 26-53, Interface-Desc.
- **l2c-downstream-data**—Exclude Juniper Networks VSA 26-93, L2C-Down-Stream-Data.
- **l2c-upstream-data**—Exclude Juniper Networks VSA 26-92, L2C-Up-Stream-Data.
- **l2tp-rx-connect-speed**—Exclude Juniper Networks VSA 26-163, Rx-Connect-Speed.
- **l2tp-tx-connect-speed**—Exclude Juniper Networks VSA 26-162, Tx-Connect-Speed.
- **max-data-rate-dn**—Exclude Juniper Networks VSA 26-120, Max-Data-Rate-Dn.

- **max-data-rate-up**—Exclude Juniper Networks VSA 26-119, Max-Data-Rate-Up.
- **max-interlv-delay-dn**—Exclude Juniper Networks VSA 26-125, Max-Interlv-Delay-Dn.
- **max-interlv-delay-up**—Exclude Juniper Networks VSA 26-123, Max-Interlv-Delay-Up.
- **min-data-rate-dn**—Exclude Juniper Networks VSA 26-116, Min-Data-Rate-Dn.
- **min-data-rate-up**—Exclude Juniper Networks VSA 26-115, Min-Data-Rate-Up.
- **min-lp-data-rate-dn**—Exclude Juniper Networks VSA 26-122, Min-Lp-Data-Rate-Dn.
- **min-lp-data-rate-up**—Exclude Juniper Networks VSA 26-121, Min-Lp-Data-Rate-Up.
- **nas-identifier**—Exclude RADIUS attribute 32, NAS-Identifier.
- **nas-port**—Exclude RADIUS attribute 5, NAS-Port.
- **nas-port-id**—Exclude RADIUS attribute 87, NAS-Port-Id.
- **nas-port-type**—Exclude RADIUS attribute 61, NAS-Port-Type.
- **output-filter**—Exclude Juniper Networks VSA 26-11, Egress-Policy-Name.
- **output-gigapackets**—Exclude Juniper Networks VSA 26-43, Acct-Output-Gigapackets.
- **output-gigawords**—Exclude RADIUS attribute 53, Acct-Output-Gigawords.
- **output-ipv6-gigawords**—Exclude Juniper Networks VSA 26-156, Acct-Output-IPv6-Gigawords.
- **output-ipv6-octets**—Exclude Juniper Networks VSA 26-152, Acct-Output-IPv6-Octets.
- **output-ipv6-packets**—Exclude Juniper Networks VSA 26-154, Acct-Output-IPv6-Packets.
- **packet-type**—Specify the RADIUS message type to exclude; term required when excluding a standard attribute or VSA by number rather than name. You can enclose multiple values in square brackets to specify a list of message types. Message types include Access-Request, Accounting-Off, Accounting-Off, Accounting-Start, and Accounting-Stop.
- **pppoe-description**—Exclude Juniper Networks VSA 26-24, PPPoE-Description.
- **standard-attribute *number***—RADIUS standard attribute number supported by your platform. If you configure an unsupported attribute, that configuration has no effect. When you use this option, you must use the **packet-type** term to specify the message from which the attribute is excluded.
- **tunnel-assignment-id**—Exclude RADIUS attribute 82, Tunnel-Assignment-ID.
- **tunnel-client-auth-id**—Exclude RADIUS attribute 90, Tunnel-Client-Auth-ID.
- **tunnel-client-endpoint**—Exclude RADIUS attribute 66, Tunnel-Client-Endpoint.
- **tunnel-medium-type**—Exclude RADIUS attribute 65, Tunnel-Medium-Type.
- **tunnel-server-auth-id**—Exclude RADIUS attribute 91, Tunnel-Server-Auth-ID.
- **tunnel-server-endpoint**—Exclude RADIUS attribute 67, Tunnel-Server-Endpoint.
- **tunnel-type**—Exclude RADIUS attribute 64, Tunnel-Type.

- **upstream-calculated-qos-rate**—Exclude Juniper Networks VSA 26-142
- **vendor-attribute** *vsa-number*—Number identifying a VSA belonging to the specified vendor; both must be supported by your platform. If you configure an unsupported VSA, that configuration has no effect. When you use this option, you must use the **packet-type** term to specify the message from which the attribute is excluded.
- **vendor-id** *id-number*—IANA vendor ID supported by your platform. If you configure an unsupported vendor ID, that configuration has no effect.
- **virtual-router**—Exclude Juniper Networks VSA 26-1.

RADIUS message type:

- **access-request**—RADIUS Access-Request messages.
- **accounting-off**—RADIUS Accounting-Off messages.
- **accounting-on**—RADIUS Accounting-On messages.
- **accounting-start**—RADIUS Accounting-Start messages.
- **accounting-stop**—RADIUS Accounting-Stop messages.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Filtering RADIUS Attributes and VSAs from RADIUS Messages

RADIUS Servers and Parameters for Subscriber Access

Standard and Vendor-Specific RADIUS Attributes

family (Address-Assignment Pools)

Syntax

```
family family {
    dhcp-attributes {
        [protocol-specific attributes]
    }
    excluded-address ip-address;
    excluded-range name low minimum-value high maximum-value;
    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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

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

Address-Assignment Pool Configuration Overview

family (Dynamic Demux Interface)

Syntax

```
family family {
  access-concentrator name;
  address address;
  demux-source {
    source-address;
  }
  direct-connect;
  duplicate-protection;
  dynamic-profile profile-name;
  filter {
    input filter-name;
    output filter-name;
  }
  mac-validate (loose | strict);
  max-sessions number;
  max-sessions-vsa-ignore;
  service-name-table table-name;
  short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max maximum-seconds> <filter [aci]>;
  unnumbered-address interface-name <preferred-source-address address>;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number]
```

Release Information

Statement introduced in Junos OS Release 9.3.

pppoe option added in Junos OS Release 11.2.

Description

Configure protocol family information for the logical interface.

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

Options

family—Protocol family:

- **inet**—Internet Protocol version 4 suite
- **inet6**—Internet Protocol version 6 suite
- **pppoe**—(MX Series routers with MPCs only) Point-to-Point Protocol over Ethernet

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles

Subscriber Interfaces and Demultiplexing Overview

family (Dynamic PPPoE)

Syntax

```
family family {
  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;
    }
    output filter-name {
      precedence precedence;
    }
  }
}
```

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 protocol family information for the logical interface.

Options

family—Protocol family:

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

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

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring a PPPoE Dynamic Profile

Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview

family (Dynamic Standard Interface)

Syntax

```
family family {
    access-concentrator name;
    address address;
    direct-connect;
    duplicate-protection;
    dynamic-profile profile-name;
    filter {
        adf {
            counter;
            input-precedence precedence;
            not-mandatory;
            output-precedence precedence;
            rule rule-value;
        }
        input filter-name {
            precedence precedence;
            shared-name filter-shared-name;
        }
        output filter-name {
            precedence precedence;
            shared-name filter-shared-name;
        }
    }
    mac-validate (loose | strict);
    max-sessions number;
    max-sessions-vsa-ignore;
    rpf-check {
        fail-filter filter-name;
        mode loose;
    }
    service {
        input {
            service-set service-set-name {
                service-filter filter-name;
            }
            post-service-filter filter-name;
        }
        output {
            service-set service-set-name {
                service-filter filter-name;
            }
        }
    }
}
```

```

    }
  }
  service-name-table table-name;
  short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max maximum-seconds> <filter [aci]>;
  unnumbered-address interface-name <preferred-source-address address>;
}

```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 9.2.

pppoe option added in Junos OS Release 11.2.

Description

Configure protocol family information for the logical interface.

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

Options

family—Protocol family:

- **inet**—IP version 4 suite
- **inet6**—IP version 6 suite
- **pppoe**—(MX Series routers with MPCs only) Point-to-Point Protocol over Ethernet
- **vpls**—Virtual private LAN service

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Static Routing on Logical Systems

Configuring the Protocol Family

fields (Flat-File Accounting Options)

Syntax

```
fields {  
  all-fields;  
  egress-stats {  
    all-fields;  
    input-bytes;  
    input-packets;  
    output-bytes;  
    output-packets;  
    queue-id;  
    red-drop-bytes;  
    red-drop-packets;  
    tail-drop-packets;  
  }  
  general-param {  
    all-fields;  
    accounting-type;  
    descr;  
    line-id;  
    logical-interface;  
    nas-port-id;  
    physical-interface;  
    routing-instance;  
    timestamp;  
    user-name;  
    vlan-id;  
  }  
  ingress-stats {  
    all-fields;  
    drop-packets;  
    input-bytes;  
    input-packets;  
    output-bytes;  
    output-packets;  
    queue-id;  
  }  
  l2-stats {  
    all-fields;  
    input-mcast-bytes;  
    input-mcast-packets;  
  }  
  overall-packet {
```

```

    all-fields;
    input-bytes;
    input-discards;
    input-errors;
    input-packets;
    inputv6-bytes;
    inputv6-packets;
    output-bytes;
    output-errors;
    output-packets;
    outputv6-bytes;
    outputv6-packets;
  }
  service-accounting;
}

```

Hierarchy Level

```
[edit accounting-options flat-file-profile profile-name]
```

Release Information

Statement introduced in Junos OS Release 16.1R4.

Description

Specify the accounting statistics and the nonstatistical information to be collected for an interface and recorded in the accounting flat file created by the profile.

Options

all-fields—Include all available statistical fields in the accounting file. Many fields are available for both logical interfaces and physical interfaces, but some fields are available only for one or the other interface type.

service-accounting—Include the filter counts in bytes for the inet input filter, inet output filter, inet6 input filter, and inet6 output filter in the service accounting flat file. Statistics reported are the running total values.

NOTE: Starting in Junos OS Release 18.4R1, the **service-accounting** option is no longer supported. If included in a configuration, it is ignored.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Flat-File Accounting for Layer 2 Wholesale 153
Configuring Flat-File Accounting for Extensible Subscriber Services Management 158
Configuring Service Accounting in Local Flat Files 163
Flat-File Accounting Overview 149

file (Flat-File Accounting Options)

Syntax

```
file filename;
```

Hierarchy Level

```
[edit accounting-options flat-file-profile profile-name]
```

Release Information

Statement introduced in Junos OS Release 16.1R4.

Description

Specify the name of the accounting file created by a flat-file profile. By default, the filename becomes the name of the local directory where the accounting file is backed up: `/var/log/pfedBackup/filename`.

Options

filename—Name of the accounting file. The complete output filename is in the format `filename.hostname.file-number_timestamp.gz`.

Required Privilege Level

system—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Flat-File Accounting for Layer 2 Wholesale | 153](#)

[Configuring Flat-File Accounting for Extensible Subscriber Services Management | 158](#)

[Flat-File Accounting Overview | 149](#)

flat-file-profile (Accounting Options)

Syntax

```
flat-file-profile profile-name{
  fields {
    all-fields;
    egress-stats {
      all-fields;
      input-bytes;
      input-packets;
      output-bytes;
      output-packets;
      queue-id;
      red-drop-bytes;
      red-drop-packets;
      tail-drop-packets;
    }
    general-param {
      all-fields;
      accounting-type;
      descr;
      line-id;
      logical-interface;
      nas-port-id;
      physical-interface;
      routing-instance;
      timestamp;
      user-name;
      vlan-id;
    }
    ingress-stats {
      all-fields;
      drop-packets;
      input-bytes;
      input-packets;
      output-bytes;
      output-packets;
      queue-id;
    }
    l2-stats {
      all-fields;
      input-mcast-bytes;
      input-mcast-packets;
    }
  }
}
```



```

overall-packet {
    all-fields;
    input-bytes;
    input-discards;
    input-errors;
    input-packets;
    inputv6-bytes;
    inputv6-packets;
    output-bytes;
    output-errors;
    output-packets;
    outputv6-bytes;
    outputv6-packets;
}
service-accounting;
}
file filename;
format (csv | ipdr)
interval minutes;
schema-version schema-name;
}

```

Hierarchy Level

[edit *accounting-options*]

Release Information

Statement introduced in Junos OS Release 16.1R4.

service-accounting option added in Junos OS Release 17.1.

Support for the **service-accounting** option removed in Junos OS Release 18.1R4.

Description

Configure a flat-file accounting profile that defines the contents of a flat file that records accounting statistics collected from the Packet Forwarding Engine for an interface at regular intervals. To be used, the profile is associated with a subscriber interface. The accounting flat file is archived by the accounting-options archiving mechanism.

Options

profile-name—Name of the flat-file profile.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Accounting-Data Log Files

[Configuring Flat-File Accounting for Layer 2 Wholesale | 153](#)

[Configuring Flat-File Accounting for Extensible Subscriber Services Management | 158](#)

[Configuring Service Accounting in Local Flat Files | 163](#)

[Flat-File Accounting Overview | 149](#)

flat-file-profile (Extensible Subscriber Services)**Syntax**

flat-file-profile *profile-name*

Hierarchy Level

[edit system services extensible-subscriber-services]

Release Information

Statement introduced in Junos OS Release 16.1.

Description

Specify the name of an accounting flat file profile that applies to an ESSM subscriber.

Options

profile name—Name of an accounting flat file profile configured at the [edit accounting-options] hierarchy level.

Required Privilege Level

system—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Flat-File Accounting for Extensible Subscriber Services Management | 158](#)

flexible-vlan-tagging

Syntax

```
flexible-vlan-tagging;
```

Hierarchy Level

```
[edit interfaces aex],
[edit interfaces ge-fpc/pic/port],
[edit interfaces et-fpc/pic/port],
[edit interfaces ps0],
[edit interfaces xe-fpc/pic/port]
```

Release Information

Statement introduced in Junos OS Release 8.1.

Support for aggregated Ethernet added in Junos OS Release 9.0.

Statement introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers.

Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.

Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.

Description

Support simultaneous transmission of 802.1Q VLAN single-tag and dual-tag frames on logical interfaces on the same Ethernet port, and on pseudowire logical interfaces.

This statement is supported on M Series and T Series routers, for Fast Ethernet and Gigabit Ethernet interfaces only on Gigabit Ethernet IQ2 and IQ2-E, IQ, and IQE PICs, and for aggregated Ethernet interfaces with member links in IQ2, IQ2-E, and IQ PICs or in MX Series DPCs, or on Ethernet interfaces for PTX Series Packet Transport Routers or 100-Gigabit Ethernet Type 5 PIC with CFP.

This statement is supported on Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces on EX Series and QFX Series switches.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Enabling VLAN Tagging

Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers

Configuring Double-Tagged VLANs on Layer 3 Logical Interfaces

format (Flat-File Accounting Options)

Syntax

```
format (csv | ipdr);
```

Hierarchy Level

```
[edit accounting-options flat-file-profile profile-name]
```

Release Information

Statement introduced in Junos OS Release 16.1R4.

Description

Specify the format for logging the flat-file accounting statistics.

Options

csv—Comma-separated values (CSV) format.

ipdr—IP Detail Record (IPDR) format.

Default: ipdr

Required Privilege Level

system—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Flat-File Accounting for Layer 2 Wholesale](#) | 153

[Configuring Flat-File Accounting for Extensible Subscriber Services Management](#) | 158

[Flat-File Accounting Overview](#) | 149

forwarding-options

Syntax

```
forwarding-options {
  access-security ...
  accounting name ...
  dhcp-relay ...
  enhanced-hash-key ...
  family family-name ...
  fast-reroute-priority (high | low | medium);
  helpers ...
  ip-options-protocol-queue protocol-name ...
  link-layer-broadcast-inet-check;
  load-balance ...
  load-balance-label-capability;
  multicast-replication ...
  next-hop-group group-name ...
  no-load-balance-label-capability;
  port-mirroring ...
  rpf-loose-mode-discard ...
  sampling ...
  satellite ...
  storm-control-profiles profile-name ...
}
```

Hierarchy Level

```
[edit]
[edit routing-instance routing-instance-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure traffic forwarding options. For more information about each option, click a linked statement in the Syntax section.

Options

access-security—Configure IPv6 access security options.

accounting—Specify the discard accounting instance name and options.

dhcp-relay—Configure extended Dynamic Host Configuration Protocol (DHCP) relay and DHCPv6 relay options on the router or switch to enable the router (or switch) to function as a DHCP relay agent.

enhanced-hash-key—(MX Series routers with MPCs, T4000 routers with Type 5 FPCs, and EX9200 switches) Select data used in the hash key for enhanced IP forwarding engines.

family—Specify address family for filters.

fast-reroute-priority—Specify the fast reroute priority for a VPLS routing instance. You can configure **high**, **medium**, or **low** fast reroute priority to prioritize specific VPLS routing instances for faster convergence and traffic restoration.

helpers—Enable TFTP or DNS request packet forwarding, or configure the router, switch, or interface to act as a DHCP/BOOTP relay agent. Enable forwarding LAN broadcast traffic on custom UDP ports to particular servers as unicast traffic.

ip-options-protocol-queue—Configure logical queue-depth in the PFE for ip-options packets for a given protocol such as TCP, UDP, ICMP, and so on, except IGMP.

link-layer-broadcast-inet-check—Enable destination MAC and IP address check to prevent the router from forwarding IPv4 packets that have link layer destination address set to **broadcast** or **multicast**, unless directed to an IPv4 multicast address.

load-balance—Enable per-prefix or per-flow load balancing so that the router or switch elects a next hop independently of the route selected by other routers or switches.

load-balance-label-capability—Enable the router to push and pop the load balancing label and causes LDP and RSVP to advertise the entropy label TLV to neighboring routers.

multicast-replication—Configure the mode of multicast replication that helps to optimize multicast latency.

next-hop-group—Specify the next-hop address for sending copies of packets to an analyzer.

no-load-balance-label-capability—Disable advertisement of entropy label capability in LDP and RSVP.

port-mirroring—Specify the input, output, and traceoptions properties for sending copies of packets to an analyzer.

rpf-loose-mode-discard—Configure unicast reverse path forwarding (unicast RPF) loose mode with the ability to discard packets with the source address pointing to the discard next hop.

sampling—Configure traffic sampling.

satellite—Specify that mirrored traffic is to be forwarded on all satellite devices that are connected to the aggregation device.

storm-control-profiles—Configure a storm control profile on a switch or router. Storm control is used to prevent network outages that are caused by broadcast traffic storms.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Stateless IPv6 Router Advertisement Guard

Configuring Stateful IPv6 Router Advertisement Guard

Configuring Discard Accounting

Extended DHCP Relay Agent Overview

Configuring Symmetrical Load Balancing on an 802.3ad Link Aggregation Group on MX Series Routers

Applying Forwarding Table Filters

Configuring VPLS Fast Reroute Priority

Configuring DNS and TFTP Packet Forwarding

Configuring Port-based LAN Broadcast Packet Forwarding

Configuring Per-Flow Load Balancing Based on Hash Values

Configuring Per-Prefix Load Balancing

Configuring the Entropy Label for LSPs

Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers

Understanding Unicast RPF (Routers)

Configuring Traffic Sampling on MX, M and T Series Routers

Understanding Port Mirroring on a Junos Fusion Data Center

OBSOLETE: Understanding Storm Control for Managing Traffic Levels

general-param (Flat-File Accounting Options)

Syntax

```
general-param {
  all-fields;
  accounting-type;
  descr;
  line-id;
  logical-interface;
  nas-port-id;
  physical-interface;
  routing-instance;
  timestamp;
  user-name;
  vlan-id;
}
```

Hierarchy Level

[edit accounting-options flat-file-profile *profile-name* **fields**]

Release Information

Statement introduced in Junos OS Release 16.1R4.

user-name option added in Junos OS Release 17.1.

Description

Specify general, nonstatistical interface parameters that are displayed as part of the header for the accounting file.

Options

all-fields—Display all available nonstatistical fields. Many fields are available for both logical interfaces and physical interfaces, but some fields are available for only one interface type.

accounting-type—(Logical interfaces only) Display the accounting status type.

descr—Display the description of the interface as configured.

line-id—(Logical interfaces only) Display the access line identifier.

logical-interface—(Logical interfaces only) Display the name of the logical interface.

nas-port-id—(Logical interfaces only) Display the NAS port ID.

physical-interface—(Physical interfaces only) Display the name of the physical interface.

routing-instance—Display the name of the routing instance to which the interface belongs.

timestamp—Display the timestamp of the accounting record.

user-name—Display the name of the subscriber.

vlan-id—(Logical interfaces only) Display the VLAN identifier.

Required Privilege Level

system—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Flat-File Accounting for Layer 2 Wholesale | 153](#)

[Configuring Flat-File Accounting for Extensible Subscriber Services Management | 158](#)

[Configuring Service Accounting in Local Flat Files | 163](#)

[Flat-File Accounting Overview | 149](#)

group (DHCP Local Server)

Syntax

```
group group-name {
  access-profile profile-name;
  authentication {
    password password-string;
    username-include {
      circuit-type;
      client-id;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-description (device-interface | logical-interface);
      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;
      vlan-tags;
    }
  }
  dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary primary-profile-name>;
  interface interface-name {
    access-profile profile-name;
    exclude;
    overrides {
      asymmetric-lease-time seconds;
      asymmetric-prefix-lease-time seconds;
      client-discover-match <option60-and-option82>;
      client-negotiation-match incoming-interface;
      delay-advertise {
        based-on (option-15 | option-16 | option-18 | option-37) {
          equals {
            ascii ascii-string;
            hexadecimal hexadecimal-string;
          }
          not-equals {
            ascii ascii-string;
            hexadecimal hexadecimal-string;
          }
        }
      }
    }
  }
}
```

```

        starts-with {
            ascii ascii-string;
            hexadecimal hexadecimal-string;
        }
    }
    delay-time seconds;
}
delay-offer {
    based-on (option-60 | option-77 | option-82) {
        equals {
            ascii ascii-string;
            hexadecimal hexadecimal-string;
        }
        not-equals {
            ascii ascii-string;
            hexadecimal hexadecimal-string;
        }
        starts-with {
            ascii ascii-string;
            hexadecimal hexadecimal-string;
        }
    }
    delay-time seconds;
}
dual-stack dual-stack-group-name;
interface-client-limit number;
process-inform {
    pool pool-name;
}
rapid-commit;
}
service-profile dynamic-profile-name;
short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
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;  
    }  
    layer2-liveness-detection {  
      max-consecutive-retries number;  
      transmit-interval interval;  
    }  
  }  
}
```

```

overrides {
  asymmetric-lease-time seconds;
  asymmetric-prefix-lease-time seconds;
  client-discover-match <option60-and-option82>;
  client-negotiation-match incoming-interface;
  delay-advertise {
    based-on (option-15 | option-16 | option-18 | option-37) {
      equals {
        ascii ascii-string;
        hexadecimal hexadecimal-string;
      }
      not-equals {
        ascii ascii-string;
        hexadecimal hexadecimal-string;
      }
      starts-with {
        ascii ascii-string;
        hexadecimal hexadecimal-string;
      }
    }
    delay-time seconds;
  }
  delay-offer {
    based-on (option-60 | option-77 | option-82) {
      equals {
        ascii ascii-string;
        hexadecimal hexadecimal-string;
      }
      not-equals {
        ascii ascii-string;
        hexadecimal hexadecimal-string;
      }
      starts-with {
        ascii ascii-string;
        hexadecimal hexadecimal-string;
      }
    }
    delay-time seconds;
  }
  delegated-pool;
  delete-binding-on-renegotiation;
  dual-stack dual-stack-group-name;
  interface-client-limit number;
  process-inform {

```

```

    pool pool-name;
  }
  protocol-attributes attribute-set-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;
short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}

```

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

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

RELATED DOCUMENTATION

<i>Understanding Differences Between Legacy DHCP and Extended DHCP</i>
<i>Grouping Interfaces with Common DHCP Configurations</i>
<i>Specifying Authentication Support</i>
<i>Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces</i>
<i>DHCP Liveness Detection Using ARP and Neighbor Discovery Packets</i>

group (DHCP Relay Agent)

Syntax

```
group group-name {
  access-profile profile-name;
  active-server-group server-group-name;
  authentication {
    password password-string;
    username-include {
      circuit-type;
      client-id;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-description (device-interface | logical-interface);
      interface-name interface-name;
      logical-system-name;
      mac-address mac-address;
      relay-agent-interface-id;
      relay-agent-remote-id;
      relay-agent-subscriber-id;
      routing-instance-name;
      user-prefix user-prefix-string;
      vlan-tags;
    }
  }
  dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
  }
  forward-only {
    logical-system <current | default | logical-system-name>;
    routing-instance <current | default | routing-instance-name>;
  }
  interface interface-name {
    access-profile profile-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 {
    allow-no-end-option;
    allow-snooped-clients;
    always-write-giaddr;
    always-write-option-82;
    asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    client-discover-match <option60-and-option82 | incoming-interface>;
    client-negotiation-match incoming-interface;
    delay-authentication;
    delete-binding-on-renegotiation;
    disable-relay;
    dual-stack dual-stack-group-name;
    interface-client-limit number;
    layer2-unicast-replies;
    no-allow-snooped-clients;
    no-bind-on-request;
    proxy-mode;
    relay-source
    replace-ip-source-with;
    send-release-on-delete;
    trust-option-82;
}
service-profile dynamic-profile-name;
short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
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;  
    }  
    layer2-liveness-detection {  
      max-consecutive-retries number;  
      transmit-interval interval;  
    }  
  }  
}
```

```

overrides {
    allow-snooped-clients;
    always-write-giaddr;
    always-write-option-82;
    asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    client-discover-match <option60-and-option82>;
    client-negotiation-match incoming-interface;
    disable-relay;
    dual-stack dual-stack-group-name;
    interface-client-limit number;
    layer2-unicast-replies;
    no-allow-snooped-clients;
    no-bind-on-request;
    proxy-mode;
    relay-source
    replace-ip-source-with;
    send-release-on-delete;
    trust-option-82;
}
relay-agent-interface-id {
    include-irb-and-l2;
    keep-incoming-interface-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82 <strict>;
    use-vlan-id;
}
relay-agent-remote-id {
    include-irb-and-l2;
    keep-incoming-remote-id ;
    no-vlan-interface-name;
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82 <strict>;
    use-vlan-id;
}

```

```

relay-option {
  option-number option-number;
  default-action {
    drop;
    forward-only;
    local-server-group local-server-group;
    relay-server-group relay-server-group;
  }
  equals (ascii ascii-string | hexadecimal hexadecimal-string) {
    drop;
    forward-only;
    local-server-group local-server-group;
    relay-server-group relay-server-group;
  }
  starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
    drop;
    forward-only;
    local-server-group local-server-group;
    relay-server-group relay-server-group;
  }
}
relay-option-82 {
  circuit-id {
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82;
  }
  remote-id {
    prefix prefix;
    use-interface-description (logical | device);
  }
  server-id-override
}
remote-id-mismatch disconnect;
route-suppression;
service-profile dynamic-profile-name;
short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
}

```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Release Information

Statement introduced in Junos OS Release 8.3.

Support at the **[edit ... dhcpv6]** hierarchy levels introduced in Junos OS Release 11.4.

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

Description

Specify the name of a group of interfaces that have a common DHCP or DHCPv6 relay agent configuration. A group must contain at least one interface. Use the statement at the **[edit ... dhcpv6]** hierarchy levels to configure DHCPv6 support.

Options

group-name—Name of a group of interfaces that have a common DHCP or DHCPv6 relay agent configuration.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Extended DHCP Relay Agent Overview](#)

[Configuring DHCP Relay Agent](#)

[Configuring Group-Specific DHCP Relay Options](#)

[Grouping Interfaces with Common DHCP Configurations](#)

[Specifying Authentication Support](#)

[Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces](#)

[DHCP Liveness Detection Using ARP and Neighbor Discovery Packets](#)

ingress-stats (Flat-File Accounting Options)

Syntax

```
ingress-stats {  
    all-fields;  
    drop-packets;  
    input-bytes;  
    input-packets;  
    output-bytes;  
    output-packets;  
    queue-id;  
}
```

Hierarchy Level

[edit accounting-options flat-file-profile *profile-name* [fields](#)]

Release Information

Statement introduced in Junos OS Release 16.1R4.

Description

Specify ingress queue statistics to be collected for the interface.

Options

all-fields—Collect all ingress queue statistics available for the interface context, logical or physical.

drop-packets—Collect the number of received packets dropped on the Ingress queue.

input-bytes—Collect the number of octets received on the queue for the traffic class indicated by the queue identifier.

input-packets—Collect the number of packets received on the queue for the traffic class indicated by the queue identifier.

output-bytes—Collect the number of octets forwarded for the traffic class indicated by the queue identifier. Same value as **input-bytes** unless oversubscription is present at the ingress.

output-packets—Collect the number of packets forwarded for the traffic class indicated by the queue identifier. Same value as **input-packets** unless oversubscription is present at the ingress.

queue-id—Collect the logical identifier for the ingress queue; identifies the traffic class.

Required Privilege Level

system—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Flat-File Accounting for Layer 2 Wholesale | 153](#)

[Configuring Flat-File Accounting for Extensible Subscriber Services Management | 158](#)

[Flat-File Accounting Overview | 149](#)

inner-vlan-id (Dynamic VLANs)

Syntax

```
inner-vlan-id number;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number input-vlan-map],  
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number output-vlan-map]
```

Release Information

Statement introduced in Junos OS Release 10.4.

Description

For dynamic VLAN interfaces, specify the VLAN ID to rewrite for the inner tag of the final packet.

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

Options

number—VLAN ID number. When used for input VLAN maps, you can specify the **\$junos-inner-vlan-map-id** predefined variable to dynamically obtain the VLAN identifier.

Range: 0 through 4094

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| *Configuring Inner and Outer TPIDs and VLAN IDs*

inner-vlan-id-swap-ranges

Syntax

```
inner-vlan-id-swap-ranges low-inner-tag-high-inner-tag;
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 17.2R1.

Description

Define core-facing VLAN ID ranges from which an inner VLAN ID tag can be allocated to replace the outer VLAN tag that was appended by the access node on the upstream packets to the BNG in a Layer 2 wholesale network. The tag swap occurs before the subscriber traffic is forwarded to the network service provider (NSP). You can configure up to 32 non-overlapping inner VLAN ID ranges per core-facing physical interface for VLAN-OOB subscribers.

You can add or remove ranges or increase or decrease the size of existing ranges even while Layer 2 wholesale sessions are assigned to the core-facing interface associated with the ranges. You cannot remove a range from which a VLAN ID has already been allocated. You cannot reduce a range if the new range excludes a VLAN ID that has already been allocated.

Options

low-inner-tag—Inner (core-facing) VLAN ID tag representing the lower limit of the swap range.

Range: 1 through 4094.

high-inner-tag—Inner (core-facing) VLAN ID tag representing the upper limit of the swap range.

Range: 1 through 4094.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Multiple Non-Overlapping VLAN Ranges for Core-Facing Physical Interfaces](#) | 146

[Layer 2 Wholesale with ANCP-Triggered VLANs Overview](#) | 116

input-vlan-map (Dynamic Interfaces)

Syntax

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

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 10.4.

Description

For dynamic interfaces, define the rewrite profile to be applied to incoming frames on this logical interface.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution](#) | 94

instance-role

Syntax

```
instance-role (access | nni);
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 11.2.

Description

Define the role of the routing instance in a Layer 2 Wholesale network.

Options

access—Defines the connectivity role of the routing instance in a Layer 2 Wholesale network as an access routing instance. When defined for this role, the same process occurs as in a Layer 3 Wholesale network—when the first packet is received from a given client, authentication for the client initiates with an external entity (for example, RADIUS). If authentication is successful, a logical interface is created with the appropriate outer and inner VLAN tags for that client.

nni—Defines the connectivity role of the routing instance in a Layer 2 Wholesale network as a network to network interface (NNI) routing instance. When defined for this role, only outer VLAN tags are learned. In addition, when the NNI routing instance receives a response from the ISP, the packets are forwarded to the appropriate client, provided the packet has the same two tags that were verified during authentication.

NOTE: If you connect an access node or MSAN device to a router participating in the Layer 2 Wholesale network in an NNI role, you must create a new routing instance of type **l2backhaul-vpn** with an instance role of type **access** for that connection.

Required Privilege Level

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Separate Access Routing Instances for Layer 2 Wholesale Service Retailers | 102](#)

[Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers | 105](#)

Subscriber Management Overview

instance-type

Syntax

```
instance-type type;
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

virtual-switch and **layer2-control** options introduced in Junos OS Release 8.4.

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

Statement introduced in Junos OS Release 11.3 for the QFX Series.

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

mpls-internet-multicast option introduced in Junos OS Release 11.1 for the EX Series, M Series, MX Series, and T Series.

evpn option introduced in Junos OS Release 13.2 for MX 3D Series routers.

evpn option introduced in Junos OS Release 17.3 for the QFX Series.

forwarding option introduced in Junos OS Release 14.2 for the PTX Series.

mpls-forwarding option introduced in Junos OS Release 16.1 for the MX Series.

evpn-vpws option introduced in Junos OS Release 17.1 for MX Series routers.

Support for logical systems on MX Series routers added in Junos OS Release 17.4R1.

Description

Define the type of routing instance.

Options

NOTE: On ACX Series routers, you can configure only the forwarding, virtual router, and VRF routing instances.

type—Can be one of the following:

- **evpn**—(MX 3D Series routers, QFX switches and EX9200 switches)—Enable an Ethernet VPN (EVPN) on the routing instance.
hierarchy level.
- **evpn-vpws**—Enable an Ethernet VPN (EVPN) Virtual Private Wire Service (VPWS) on the routing instance.

- **forwarding**—Provide support for filter-based forwarding, where interfaces are not associated with instances. All interfaces belong to the default instance. Other instances are used for populating RPD learned routes. For this instance type, there is no one-to-one mapping between an interface and a routing instance. All interfaces belong to the default instance inet.0.
- **l2backhaul-vpn**—Provide support for Layer 2 wholesale VLAN packets with no existing corresponding logical interface. When using this instance, the router learns both the outer tag and inner tag of the incoming packets, when the **instance-role** statement is defined as **access**, or the outer VLAN tag only, when the **instance-role** statement is defined as **nni**.
- **l2vpn**—Enable a Layer 2 VPN on the routing instance. You must configure the **interface**, **route-distinguisher**, **vrf-import**, and **vrf-export** statements for this type of routing instance.
- **layer2-control**—(MX Series routers only) Provide support for RSTP or MSTP in customer edge interfaces of a VPLS routing instance. This instance type cannot be used if the customer edge interface is multihomed to two provider edge interfaces. If the customer edge interface is multihomed to two provider edge interfaces, use the default BPDU tunneling.
- **mpls-forwarding**—(MX Series routers only) Allow filtering and translation of route distinguisher (RD) values in IPv4 and IPv6 VPN address families on both routes received and routes sent for selected BGP sessions. In particular, for Inter-AS VPN Option-B networks, this option can prevent the malicious injection of VPN labels from one peer AS boundary router to another.
- **mpls-internet-multicast**—(EX Series, M Series, MX Series, and T Series routers only) Provide support for ingress replication provider tunnels to carry IP multicast data between routers through an MPLS cloud, using MBGP or next-generation MVPN.
- **no-forwarding**—This is the default routing instance. Do not create a corresponding forwarding instance. Use this routing instance type when a separation of routing table information is required. There is no corresponding forwarding table. All routes are installed into the default forwarding table. IS-IS instances are strictly nonforwarding instance types.
- **virtual-router**—Enable a virtual router routing instance. This instance type is similar to a VPN routing and forwarding instance type, but used for non-VPN-related applications. You must configure the **interface** statement for this type of routing instance. You do not need to configure the **route-distinguisher**, **vrf-import**, and **vrf-export** statements.
- **virtual-switch**—(MX Series routers, EX9200 switches, and QFX switches only) Provide support for Layer 2 bridging. Use this routing instance type to isolate a LAN segment with its Spanning Tree Protocol (STP) instance and to separate its VLAN identifier space.

- **vpls**—Enable VPLS on the routing instance. Use this routing instance type for point-to-multipoint LAN implementations between a set of sites in a VPN. You must configure the **interface**, **route-distinguisher**, **vrf-import**, and **vrf-export** statements for this type of routing instance.
- **vrf**—VPN routing and forwarding (VRF) instance. Provides support for Layer 3 VPNs, where interface routes for each instance go into the corresponding forwarding table only. Required to create a Layer 3 VPN. Create a VRF table (**instance-name.inet.0**) that contains the routes originating from and destined for a particular Layer 3 VPN. For this instance type, there is a one-to-one mapping between an interface and a routing instance. Each VRF instance corresponds with a forwarding table. Routes on an interface go into the corresponding forwarding table. You must configure the **interface**, **route-distinguisher**, **vrf-import**, and **vrf-export** statements for this type of routing instance.

Required Privilege Level

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring the Instance Type

Configuring EVPN Routing Instances

Configuring EVPN Routing Instances on EX9200 Switches

Configuring Virtual Router Routing Instances

Example: Configuring Filter-Based Forwarding on the Source Address

Example: Configuring Filter-Based Forwarding on Logical Systems

interface (DHCP Local Server)

Syntax

```
interface interface-name {
  access-profile profile-name;
  exclude;
  overrides {
    asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    client-discover-match <option60-and-option82 | incoming-interface>;
    client-negotiation-match incoming-interface;
    delay-advertise {
      based-on (option-15 | option-16 | option-18 | option-37) {
        equals {
          ascii ascii-string;
          hexadecimal hexadecimal-string;
        }
        not-equals {
          ascii ascii-string;
          hexadecimal hexadecimal-string;
        }
        starts-with {
          ascii ascii-string;
          hexadecimal hexadecimal-string;
        }
      }
      delay-time seconds;
    }
    delay-offer {
      based-on (option-60 | option-77 | option-82) {
        equals {
          ascii ascii-string;
          hexadecimal hexadecimal-string;
        }
        not-equals {
          ascii ascii-string;
          hexadecimal hexadecimal-string;
        }
        starts-with {
          ascii ascii-string;
          hexadecimal hexadecimal-string;
        }
      }
      delay-time seconds;
    }
  }
}
```



```

    }
    dual-stack dual-stack-group-name;
    interface-client-limit number;
    rapid-commit;
  }
  service-profile dynamic-profile-name;
  short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
  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 supports only static DHCP configurations.

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Understanding Differences Between Legacy DHCP and Extended DHCP

Grouping Interfaces with Common DHCP Configurations

Specifying Authentication Support

interface (DHCP Relay Agent)

Syntax

```
interface dhcp-interface-name {
  access-profile profile-name;
  exclude;
  overrides {
    allow-no-end-option
    allow-snooped-clients;
    always-write-giaddr;
    always-write-option-82;
    asymmetric-lease-time seconds;
    asymmetric-prefix-lease-time seconds;
    client-discover-match <option60-and-option82 | incoming-interface>;
    client-negotiation-match incoming-interface;
    disable-relay;
    dual-stack dual-stack-group-name;
    interface-client-limit number;
    layer2-unicast-replies;
    no-allow-snooped-clients;
    proxy-mode;
    relay-source
    replace-ip-source-with;
    send-release-on-delete;
    trust-option-82;
  }
  service-profile dynamic-profile-name;
  short-cycle-protection <lockout-min-time seconds> <lockout-max-time seconds>;
  trace;
  upto upto-interface-name;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit forwarding-options dhcp-relay group group-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Release Information

Statement introduced in Junos OS Release 8.3.

Options **upto** and **exclude** introduced in Junos OS Release 9.1.

Support at the **[edit ... dhcpv6]** hierarchy levels introduced in Junos OS Release 11.4.

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

Description

Specify one or more interfaces, or a range of interfaces, that are within a specified group on which the DHCP or DHCPv6 relay agent 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 local server. Use the statement at the **[edit ... dhcpv6]** hierarchy levels to configure DHCPv6 support.

EX Series switches do not support DHCPv6.

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. .

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.

overrides—Override the specified default configuration settings for the interface. The **overrides** statement is described separately.

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview

Grouping Interfaces with Common DHCP Configurations

Specifying Authentication Support

interface (Dynamic Routing Instances)

Syntax

```
interface interface-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name routing-instances routing-instance-name]
```

Release Information

Statement introduced in Junos OS Release 9.6.

Description

Assign the specified interface to the dynamically created routing instance.

Options

interface-name—The interface name variable (*\$junos-interface-name*). The interface name variable is dynamically replaced with the interface the accessing client uses when connecting to the router.

Required Privilege Level

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

interface (Routing Instances)

Syntax

```
interface interface-name {
    description text;
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

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

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

Statement introduced in Junos OS Release 13.2 for MX 3D Series routers.

Description

Specify the interface over which the VPN traffic travels between the PE device and CE device. You configure the interface on the PE device. If the value **vrf** is specified for the **instance-type** statement included in the routing instance configuration, this statement is required.

Options

interface-name—Name of the interface.

The remaining statement is explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Interfaces for VPN Routing](#)

[Configuring EVPN Routing Instances](#)

[Configuring EVPN Routing Instances on EX9200 Switches](#)

[interface \(VPLS Routing Instances\)](#)

interface-mac-limit (VPLS)

Syntax

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

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls site site-name
  interfaces interface-name],
[edit routing-instances routing-instance-name protocols evpn],
[edit routing-instances routing-instance-name protocols evpn interface interface-name],
[edit routing-instances routing-instance-name protocols vpls site site-name interfaces interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Support for EVPNs introduced in Junos OS Release 13.2 on MX 3D Series routers.

Support for EVPNs introduced in Junos OS Release 14.2 on EX Series switches.

Description

Specify the maximum number of media access control (MAC) addresses that can be learned by the EVPN or VPLS routing instance. You can configure the same limit for all interfaces configured for a routing instance. You can also configure a limit for a specific interface.

Starting with Junos OS Release 12.3R4, if you do not configure the parameter to limit the number of MAC addresses to be learned by a VPLS instance, the default value is not effective. Instead, if you do not include the **interface-mac-limit** option at the **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls site *site-name* interfaces *interface-name*]**, hierarchy level, this setting is not present in the configuration with the default value of 1024 addresses. If you upgrade a router running a Junos OS release earlier than Release 12.3R4 to Release 12.3R4 or later, you must configure the **interface-mac-limit** option with a valid value for it to be saved in the configuration.

Options

limit—Number of MAC addresses that can be learned from each interface.

Range: 1 through 131,071 MAC addresses

NOTE: For M120 devices only, the range is 16 through 65,536 MAC addresses.

Default: 1024 addresses

The remaining statement is explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

<i>Configuring EVPN Routing Instances</i>
<i>Configuring EVPN Routing Instances on EX9200 Switches</i>
<i>Limiting the Number of MAC Addresses Learned from an Interface</i>
<i>interface</i>
<i>mac-table-size</i>

interfaces (Static and Dynamic Subscribers)

Syntax

```

interfaces {
  interface-name {
    unit logical-unit-number {
      actual-transit-statistics;
      auto-configure {
        agent-circuit-identifier {
          dynamic-profile profile-name;
        }
        line-identity {
          include {
            accept-no-ids;
            circuit-id;
            remote-id;
          }
          dynamic-profile profile-name;
        }
      }
    }
    family family {
      access-concentrator name;
      address address;
      direct-connect;
      duplicate-protection;
      dynamic-profile profile-name;
      filter {
        adf {
          counter;
          input-precedence precedence;
          not-mandatory;
          output-precedence precedence;
          rule rule-value;
        }
        input filter-name {
          precedence precedence;
          shared-name filter-shared-name;
        }
        output filter-name {
          precedence precedence;
          shared-name filter-shared-name;
        }
      }
      max-sessions number;
    }
  }
}

```

```

max-sessions-vsa-ignore;
rpf-check {
    mode loose;
}
service {
    input {
        service-set service-set-name {
            service-filter filter-name;
        }
        post-service-filter filter-name;
    }
    output {
        service-set service-set-name {
            service-filter filter-name;
        }
    }
}
service-name-table table-name
short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max maximum-seconds>;
unnumbered-address interface-name <preferred-source-address address>;
}
filter {
    input filter-name (
        precedence precedence;
        shared-name filter-shared-name;
    )
    output filter-name {
        precedence precedence;
        shared-name filter-shared-name;
    }
}
host-prefix-only;
ppp-options {
    chap;
    pap;
}
proxy-arp;
service {
    pcef pcef-profile-name {
        activate rule-name | activate-all;
    }
}

```

```

targeted-options {
    backup backup;
    group group;
    primary primary;
    weight ($junos-interface-target-weight | weight-value);
}
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;
      direct-connect;
      duplicate-protection;
      dynamic-profile profile-name;
      demux-source {
        source-prefix;
      }
      filter {
        input filter-name (
          precedence precedence;
          shared-name filter-shared-name;
        )
        output filter-name {
          precedence precedence;
          shared-name filter-shared-name;
        }
      }
      mac-validate (loose | strict):
      max-sessions number;
      max-sessions-vsa-ignore;
      rpf-check {
        fail-filter filter-name;
        mode loose;
      }
      service-name-table table-name
      short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max maximum-seconds>;
      unnumbered-address interface-name <preferred-source-address address>;
    }
    filter {
      input filter-name;
      output filter-name;
    }
    vlan-id number;
    vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
  }
}

```

```

pp0 {
  unit logical-unit-number {
    keepalives interval seconds;
    no-keepalives;
    pppoe-options {
      underlying-interface interface-name;
      server;
    }
    ppp-options {
      aaa-options aaa-options-name;
      authentication [ authentication-protocols ];
      chap {
        challenge-length minimum minimum-length maximum maximum-length;
        local-name name;
      }
      ignore-magic-number-mismatch;
      initiate-ncp (dual-stack-passive | ipv6 | ip)
      ipcp-suggest-dns-option;
      mru size;
      mtu (size | use-lower-layer);
      on-demand-ip-address;
      pap;
      peer-ip-address-optional;
      local-authentication {
        password password;
        username-include {
          circuit-id;
          delimiter character;
          domain-name name;
          mac-address;
          remote-id;
        }
      }
    }
  }
  family inet {
    unnumbered-address interface-name;
    address address;
    service {
      input {
        service-set service-set-name {
          service-filter filter-name;
        }
        post-service-filter filter-name;
      }
    }
  }
}

```

```

        output {
            service-set service-set-name {
                service-filter filter-name;
            }
        }
    }
    filter {
        input filter-name {
            precedence precedence;
            shared-name filter-shared-name;
        }
        output filter-name {
            precedence precedence;
            shared-name filter-shared-name;
        }
    }
}
}
}
stacked-interface-set {
    interface-set-name interface-set-name {
        interface-set-name interface-set-name;
    }
}
}
}

```

Hierarchy Level

[edit [dynamic-profiles](#) *profile-name*]

Release Information

Statement introduced in Junos OS Release 9.2.

Description

Define interfaces for dynamic client 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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles

Configuring Dynamic PPPoE Subscriber Interfaces

Configuring Dynamic VLANs Based on Agent Circuit Identifier Information

DHCP Subscriber Interface Overview

Subscribers over Static Interfaces Configuration Overview

Demultiplexing Interface Overview

interval (Flat-File Accounting Options)

Syntax

```
interval minutes;
```

Hierarchy Level

```
[edit accounting-options flat-file-profile profile-name]
```

Release Information

Statement introduced in Junos OS Release 16.1R4.

Description

Specify the interval in minutes at which the Packet Forwarding Engine associated with the interface is polled to collect the statistics specified in the flat-file accounting profile. These interim accounting results are recorded in the flat file.

NOTE: The value configured with this statement is superseded by the value configured with the *update-interval* statement at the **[edit access profile *profile-name* service accounting]** hierarchy level. That access profile interval value is in turn superseded by an update interval value configured in the RADIUS attribute, Service-Interim-Acct-Interval (VSA 26–140).

Options

minutes—Polling interval.

Range: 1 through 2880 minutes

Default: 15 minutes

Required Privilege Level

system—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Flat-File Accounting for Layer 2 Wholesale | 153](#)

[Configuring Flat-File Accounting for Extensible Subscriber Services Management | 158](#)

[Configuring Service Accounting in Local Flat Files | 163](#)

[Flat-File Accounting Overview | 149](#)

ip-address-first

Syntax

```
ip-address-first;
```

Hierarchy Level

```
[edit logical-systems logical-system-name system services dhcp-local-server pool-match-order],
[edit logical-systems logical-system-name routing-instances routing-instance-name system services dhcp-local-server
pool-match-order],
[edit routing-instances routing-instance-name system services dhcp-local-server pool-match-order],
[edit system services dhcp-local-server pool-match-order]
```

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 extended DHCP local server to use the IP address method to determine which address-assignment pool to use. The local server uses the IP address in the gateway IP address if one is present in the DHCP client PDU. If no gateway IP address is present, the local server uses the IP address of the receiving interface to find the address-assignment pool. The DHCP local server uses this method by default when no method is explicitly specified.

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](#)

[Understanding Differences Between Legacy DHCP and Extended DHCP](#)

[Address-Assignment Pools Overview](#)

keepalives (Dynamic Profiles)

Syntax

```
keepalives {
    interval seconds;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces pp0 unit logical-unit-number ]
[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 pp0 unit "\$junos-interface-unit"]** hierarchy level introduced in Junos OS Release 10.1.

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

Specify the keepalive interval in a PPP dynamic profile.

Starting in Junos OS Release 15.1R5, you can configure the PPP keepalive interval for subscriber services in the range 1 second through 600 seconds. Subscriber PPP keepalives are handled by the Packet Forwarding Engine. If you configure a value greater than 600 seconds, the number is accepted by the CLI, but the Packet Forwarding Engine limits the interval to 600 seconds.

In earlier Junos OS releases, the range is from 1 second through 60 seconds. The Packet Forwarding Engine limits any higher configured value to an interval of 60 seconds.

PPP keepalives for nonsubscriber services are handled by the Routing Engine with an interval range from 1 second through 32,767 seconds.

Default

Sending of keepalives is enabled by default.

Options

interval *seconds*—The time in seconds between successive keepalive requests.

Range: 1 through 600 seconds for subscriber services

Range: 1 through 32767 seconds for nonsubscriber services

Default: 30 seconds for LNS-based PPP sessions. 10 seconds for all other PPP sessions.

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

Applying PPP Attributes to L2TP LNS Subscribers per Inline Service Interface

I2-stats (Flat-File Accounting Options)

Syntax

```
I2-stats {  
    all-fields;  
    input-mcast-bytes;  
    input-mcast-packets;  
}
```

Hierarchy Level

[edit accounting-options flat-file-profile *profile-name* [fields](#)]

Release Information

Statement introduced in Junos OS Release 16.1R4.

Description

Specify the statistics to collect for the named flat-file-profile field.

Options

all-fields—Collect all Layer 2 statistics for the named flat-file profile.

input-mcast-bytes—Collect multicast bytes from the input side for the named flat-file profile.

input-mcast-packets—Collect multicast packets from the input side for the named flat-file profile.

Required Privilege Level

system—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Flat-File Accounting for Layer 2 Wholesale](#) | 153

[Configuring Flat-File Accounting for Extensible Subscriber Services Management](#) | 158

[Flat-File Accounting Overview](#) | 149

mac-validate (Dynamic IP Demux Interface)

Syntax

```
mac-validate (loose | strict);
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family inet]
```

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Enable IP and MAC address validation for dynamic IP demux interfaces in a dynamic profile.

Options

loose—Forwards incoming packets when both the IP source address and the MAC source address match one of the trusted address tuples. Drops packets when the IP source address matches one of the trusted tuples, but the MAC address does not match the MAC address of the tuple. Continues to forward incoming packets when the source address of the incoming packet does not match any of the trusted IP addresses.

strict—Forwards incoming packets when both the IP source address and the MAC source address match one of the trusted address tuples. Drops packets when the MAC address does not match the tuple's MAC source address, or when IP source address of the incoming packet does not match any of the trusted IP addresses.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| *Configuring MAC Address Validation for Subscriber Interfaces*

multicast-replication

Syntax

```
multicast-replication {  
  evpn {  
    irb (local-only | local-remote);  
    smet-nexthop-limit smet-nexthop-limit;  
  }  
  ingress;  
  local-latency-fairness;  
}
```

Hierarchy Level

[edit forwarding-options]

Release Information

Statement introduced in Junos OS Release 15.1 for MX Series routers.

evpn stanza introduced in Junos OS Release 17.3R3 for QFX Series switches.

Description

Configure the mode of multicast replication that helps to optimize multicast latency.

NOTE: The **multicast-replication** statement is supported only on platforms with the **enhanced-ip** mode enabled.

Default

This statement is disabled by default.

Options

NOTE: The **ingress** and **local-latency-fairness** options do not apply to EVPN configurations.

ingress—Complete ingress replication of the multicast data packets where all the egress Packet Forwarding Engines receive packets from the ingress Packet Forwarding Engines directly.

local-latency-fairness—Complete parallel replication of the multicast data packets.

evpn irb local-only—Enables IPv4 inter-VLAN multicast forwarding in an EVPN-VXLAN network with a collapsed IP fabric, which is also known as a *edge-routed bridging overlay*.

evpn irb local-remote—Enables IPv4 inter-VLAN multicast forwarding in an EVPN-VXLAN network with a two-layer IP fabric, which is also known as a *centrally-routed bridging overlay*.

NOTE: Selective multicast forwarding is only supported with local-remote.

Default: `evpn irb local-remote`

smet-nexthop-limit smet-nexthop-limit—Configures a limit for the number of SMET next hops for selective multicast forwarding. SMET next hops is a list of outgoing interfaces used by a PE device in selectively replicating and forwarding multicast traffic. When this limit is reached, no new SMET next hop is created and the PE device will send the new multicast group traffic to all egress devices.

Range: 10,000 through 40,000

Default: 10,000

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[forwarding-options](#) | 345

IPv4 Inter-VLAN Multicast Forwarding Modes for EVPN-VXLAN Overlay Networks

neighbor (Define ANCP)

Syntax

```
neighbor ip-address {  
    adjacency-loss-hold-time seconds;  
    adjacency-timer;  
    auto-configure-trigger interface interface-name;  
    ietf-mode;  
    maximum-discovery-table-entries entry-number;  
    pre-ietf-mode;  
}
```

Hierarchy Level

```
[edit protocols ancp]
```

Release Information

Statement introduced in Junos OS Release 9.4.

Description

Configure an ANCP neighbor with which the ANCP agent on the router forms an adjacency for reporting and shaping traffic.

Options

ip-address—IP address of the ANCP neighbor.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring the ANCP Agent

Configuring ANCP Neighbors

no-local-switching

Syntax

```
no-local-switching
```

Hierarchy Level

```
[edit vlans vlan-name]
```

Release Information

Statement introduced in Junos OS Release 9.3 for EX Series switches.

Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Specify that access ports in this VLAN domain do not forward packets to each other. You use this statement with primary VLANs and isolated secondary VLANs.

Required Privilege Level

system—To view this statement in the configuration.

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

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Creating a Private VLAN on a Single QFX Switch

Creating a Private VLAN Spanning Multiple QFX Series Switches

no-tunnel-services

Syntax

```
no-tunnel-services;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols vpls static-vpls],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls],  
[edit protocols vpls static-vpls],  
[edit routing-instances routing-instance-name protocols vpls]
```

Release Information

Statement introduced in Junos OS Release 7.6.

Support for static VPLS added in Junos OS Release 10.2.

Description

Configure VPLS on a router without a Tunnel Services PIC. Configuring the **no-tunnel-services** statement creates a label-switched interface (LSI) to provide VPLS functionality. An LSI MPLS label is used as the inner label for VPLS. This label maps to a VPLS routing instance. On the PE router, the LSI label is stripped and then mapped to a logical LSI interface. The Layer 2 Ethernet frame is then forwarded using the LSI interface to the correct VPLS routing instance.

NOTE: In VPLS documentation, the word *Router* in terms such as *PR Router* is used to refer to any device that provides routing functions.

NOTE:

On MX Series routers, label-switched interfaces configured with the **no-tunnel-services** statement are not supported with GRE tunnels when the GRE interface resides on a DPC.

NOTE: Although visible in the CLI, the **no-tunnel-services** statement is not supported on DPC cards at the `[edit logical-systems logical-system-name protocols vpls static-vpls]` and the `[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls]` hierarchy levels.

Required Privilege Level

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring VPLS Without a Tunnel Services PIC

Configuring Static Pseudowires for VPLS

Configuring EXP-Based Traffic Classification for VPLS

overall-packet (Flat-File Accounting Options)

Syntax

```
overall-packet {  
  all-fields;  
  input-bytes;  
  input-discards;  
  input-errors;  
  input-packets;  
  inputv6-bytes;  
  inputv6-packets;  
  output-bytes;  
  output-errors;  
  output-packets;  
  outputv6-bytes;  
  outputv6-packets;  
}
```

Hierarchy Level

[edit accounting-options flat-file-profile *profile-name* [fields](#)]

Release Information

Statement introduced in Junos OS Release 16.1R4.

Description

Specify overall packet statistics to be collected for the interface.

Options

all-fields—Collect all overall packet statistics available for the interface context, logical or physical.

input-bytes—Collect the number of octets received on the interface.

input-discards—(Physical interfaces only) Collect the number of received packets that were discarded on the interface.

input-errors—(Physical interfaces only) Collect the number of frames with errors received on the interface.

input-packets—Collect the number of packets received on the interface.

input-v6-bytes—Collect the number of IPv6 octets received on the interface.

input-v6-packets—Collect the number of IPv6 packets received on the interface.

output-bytes—Collect the number of octets transmitted on the interface.

output-errors—(Physical interfaces only) Collect the number of frames that could not be transmitted on the interface because of errors.

output-packets—Collect the number of packets transmitted on the interface.

output-v6-bytes—Collect the number of IPv6 octets transmitted on the interface.

output-v6-packets—Collect the number of IPv6 packets transmitted on the interface.

Required Privilege Level

system—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Flat-File Accounting for Layer 2 Wholesale | 153](#)

[Configuring Flat-File Accounting for Extensible Subscriber Services Management | 158](#)

[Flat-File Accounting Overview | 149](#)

output-vlan-map (Dynamic Interfaces)

Syntax

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

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 10.4.

Description

For dynamic interfaces, define the rewrite profile to be applied to outgoing frames on this logical interface.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| [Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution](#) | 94

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](#)

pool (Address-Assignment Pools)

Syntax

```
pool pool-name {
  active-drain;
  family family {
    dhcp-attributes {
      [ protocol-specific attributes ]
    }
    excluded-address ip-address;
    excluded-range name low minimum-value high maximum-value;
    host hostname {
      hardware-address mac-address;
      ip-address ip-address;
    }
    network ip-prefixprefix-length>;
    prefix ipv6-prefix;
    range range-name {
      high upper-limit;
      low lower-limit;
      prefix-length prefix-length;
    }
  }
  hold-down;
  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.

active-drain—(MX Series only) Starting in Junos OS Release 17.2, configure the DHCP local server to stop allocating addresses from this pool. When this is configured, the DHCP local server gracefully shifts clients from this address pool to an alternative pool for which active drain is not configured. When existing clients with an address from this pool submit a DHCPv4 request or DHCPv6 renew, they receive a NAK, forcing them to renegotiate. The server responds with a DHCPv4 offer or DHCPv6 advertise message with an address from a different pool.

family —Configure the protocol family for the address-assignment pool.

The options for this statement are explained separately. Click the linked statement for details.

hold-down—(MX Series only) Starting in Junos OS Release 16.1, configure an address-assignment pool that is currently in use to be unavailable for further address allocation. When a pool is in the hold-down state, the pool is no longer used to allocate IP addresses for subscribers. Current subscribers who previously obtained an address from the pool are not affected; they can continue to renew their leases. As each of these users disconnects, their address is not reallocated. The pool becomes inactive when all subscribers have disconnected and their addresses are returned to the pool.

link—(M Series, MX Series, SRX Series, T Series only) Designate a secondary address-assignment pool that is linked to the pool being configured. When the pool being configured has no addresses available for allocation, the secondary pool can be searched for a free address. You can configure a chain of linked pools, but you cannot directly link more than one pool to or from any other pool. Each linked pool in the chain serves as a backup pool for the pool immediately before it in the chain.

Values: *pool-name*—Name assigned to the secondary address-assignment pool.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Address-Assignment Pools Overview](#)

[Address-Assignment Pool Configuration Overview](#)

[Configuring Address-Assignment Pool Linking](#)

[Address Allocation from Linked Address Pools](#)

[Configuring DHCP Local Address Pool Rapid Drain](#)

[Attributes That Can Be Applied to DHCP Clients](#)

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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

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
Understanding Differences Between Legacy DHCP and Extended DHCP

pop (Dynamic VLANs)

Syntax

```
pop;
```

Hierarchy Level

[edit [dynamic-profiles](#) *profile-name* [interfaces](#) *interface-name* [unit](#) *logical-unit-number* [output-vlan-map](#)]

Release Information

Statement introduced in Junos OS Release 10.4.

Description

For dynamic VLAN interfaces, specify the VLAN rewrite operation to remove a VLAN tag from the top of the VLAN tag stack. The outer VLAN tag of the frame is removed.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Removing a VLAN Tag

[Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution](#) | 94

pppoe-options (Dynamic PPPoE)

Syntax

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

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces ppO 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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring a PPPoE Dynamic Profile

Configuring Dynamic PPPoE Subscriber Interfaces

pppoe-underlying-options (Static and Dynamic Subscribers)

Syntax

```
pppoe-underlying-options {
  access-concentrator name;
  dynamic-profile profile-name;
  direct-connect
  duplicate-protection;
  max-sessions number;
  max-sessions-vsa-ignore;
  service-name-table table-name;
  short-cycle-protection <lockout-time-min minimum-seconds> <lockout-time-max maximum-seconds> <filter [aci]>;
}
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 10.0.

Description

Configure PPPoE-specific interface properties for the underlying interface on which the router creates a static or dynamic PPPoE logical interface. The underlying interface must be configured with PPPoE (**ppp-over-ether**) encapsulation.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring PPPoE (for static interfaces)

Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces

Assigning a Service Name Table to a PPPoE Underlying Interface

ppp-options (Dynamic PPP)

Syntax

```
ppp-options {
  aaa-options aaa-options-name;
  authentication [ authentication-protocols ];
  chap {
    challenge-length minimum minimum-length maximum maximum-length;
    local-name name;
  }
  ignore-magic-number-mismatch;
  initiate-ncp (dual-stack-passive | ipv6 | ip)
  ipcp-suggest-dns-option;
  lcp-connection-update;
  mru size;
  mtu (size | use-lower-layer);
  on-demand-ip-address;
  pap;
  peer-ip-address-optional;
  local-authentication {
    password password;
    username-include {
      circuit-id;
      delimiter character;
      domain-name name;
      mac-address;
      remote-id;
    }
  }
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces "$junos-interface-ifd-name" unit "$junos-interface-unit"].
[edit dynamic-profiles profile-name interfaces pp0 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.

NOTE:

PPP options can also be configured in a group profile with the **ppp-options (L2TP)** statement. The following behavior determines the interaction between the PPP options configured in a group profile and the PPP options configured in a dynamic profile:

- When PPP options are configured only in the group profile, the group profile options are applied to the subscriber.
- When PPP options are configured in both a group profile and a dynamic profile, the dynamic profile configuration takes complete precedence over the group profile when the dynamic profile includes one or more of the PPP options that can be configured in the group profile. Complete precedence means that there is no merging of options between the profiles. The group profile is applied to the subscriber only when the dynamic profile does not include any PPP option available in the group profile.

Options

lcp-connection-update—Enable PPP to act on a Connection-Status-Message VSA (26–218) received by authd in either a RADIUS Access-Accept message or a CoA message. PPP conveys the contents of the VSA in an LCP Connection-Update-Request message to the remote peer, such as a home gateway. This action requires the following to be true:

- At least the first address family has been successfully negotiated and the session is active.
- The router LCP is in the Opened state.

Otherwise PPP takes no action on the VSA. If you do not enable the **lcp-connection-update** option, PPP processes the notification from authd, but takes no action.

Default: Disabled

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Dynamic Profiles Overview

Configuring a PPPoE Dynamic Profile

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

prefix (Address-Assignment Pools)

Syntax

```
prefix ipv6-prefix;
```

Hierarchy Level

```
[edit access address-assignment pool pool-name family inet6]
```

Release Information

Statement introduced in Junos OS Release 10.0.

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

Description

Specify the IPv6 prefix for the IPv6 address-assignment pool. This statement is mandatory for IPv6 address-assignment pools.

Options

ipv6-prefix—The IPv6 prefix.

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

Address-Assignment Pool Configuration Overview

profile (Access)

Syntax

```

profile profile-name {
  accounting {
    address-change-immediate-update
    accounting-stop-on-access-deny;
    accounting-stop-on-failure;
    ancp-speed-change-immediate-update;
    coa-immediate-update;
    coa-no-override service-class-attribute;
    duplication;
    duplication-filter;
    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;
  }
  accounting-order (radius | [accounting-order-data-list]);
  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 number;
      maximum-sessions-per-tunnel number;
    }
  }
}

```

```

    multilink {
        drop-timeout milliseconds;
        fragment-threshold bytes;
    }
    override-result-code session-out-of-resource;
    ppp-authentication (chap | pap);
    ppp-profile profile-name;
    service-profile profile-name(parameter)&profile-name;
    sessions-limit-group limit-group-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;
local {
    flat-file-profile profile-name;
}
preauthentication-order preauthentication-method;
provisioning-order (gx-plus | jsr | pcrf);

```

```

radius {
  accounting-server [ ip-address ];
  attributes {
    exclude {
      attribute-name packet-type;
      standard-attribute number {
        packet-type [ access-request | accounting-off | accounting-on | accounting-start | accounting-stop ];
      }
      vendor-id id-number {
        vendor-attribute vsa-number {
          packet-type [ access-request | accounting-off | accounting-on | accounting-start | accounting-stop ];
        }
      }
    }
  }
  ignore {
    dynamic-iflset-name;
    framed-ip-netmask;
    idle-timeout;
    input-filter;
    logical-system:routing-instance;
    output-filter;
    session-timeout;
    standard-attribute number;
    vendor-id id-number {
      vendor-attribute vsa-number;
    }
  }
}
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;
    interface-text-description;
    mac-address;
    nas-identifier;
    stacked-vlan;
    vlan;
  }
  chap-challenge-in-request-authenticator;
  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-channel;
    exclude-sub-interface;
}
juniper-access-line-attributes;
nas-identifier identifier-value;
nas-port-extended-format {
    adapter-width width;
    ae-width width;
    port-width width;
    pw-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;
    interface-text-description;
    nas-identifier;
    order {
        agent-circuit-id;
        agent-remote-id;
        interface-description;
        interface-text-description;
        nas-identifier;
        postpend-vlan-tags;
    }
    postpend-vlan-tags;
}

```

```

nas-port-type {
    ethernet {
        port-type;
    }
}
override {
    calling-station-id remote-circuit-id;
    nas-ip-address tunnel-client-gateway-address;
    nas-port tunnel-client-nas-port;
    nas-port-type tunnel-client-nas-port-type;
}
remote-circuit-id-delimiter;
remote-circuit-id-fallback {
    remote-circuit-id-format;
    agent-circuit-id;
    agent-remote-id;
}
revert-interval interval;
service-activation {
    dynamic-profile (optional-at-login | required-at-login);
    extensible-service (optional-at-login | required-at-login);
}
vlan-nas-port-stacked-format;
}
preauthentication-server ip-address;
}
radius-server server-address {
    accounting-port port-number;
    accounting-retry number;
    accounting-timeout seconds;
    dynamic-request-port
    port port-number;
    preauthentication-port port-number;
    preauthentication-secret password;
    retry attempts;
    routing-instance routing-instance-name;
    secret password;
    max-outstanding-requests value;
    source-address source-address;
    timeout seconds;
}

```

```

service {
  accounting {
    statistics (time | volume-time);
    update-interval minutes;
  }
  accounting-order (activation-protocol | local | radius);
}
session-limit-per-username number;
session-options {
  client-idle-timeout minutes;
  client-idle-timeout-ingress-only;
  client-session-timeoutminutes;
  pcc-context {
    input-service-filter-name filter-name;
    input-service-set-name service-set-name;
    ipv6-input-service-filter-name filter-name;
    ipv6-input-service-set-name service-set-name;
    ipv6-output-service-filter-name filter-name;
    ipv6-output-service-set-name service-set-name;
    output-service-filter-name filter-name;
    output-service-set-name service-set-name;
    profile-name pcef-profile-name;
  }
  strip-user-name {
    delimiter [ delimiter ];
    parse-direction (left-to-right | right-to-left);
  }
}
subscriber username {
  delegated-pool delegated-pool-name;
  framed-ip-address ipv4-address;
  framed-ipv6-pool ipv6-pool-name;
  framed-pool ipv4-pool-name;
  password password;
  target-logical-system logical-system-name <target-routing-instance (default | routing-instance-name)>;
  target-routing-instance (default | routing-instance-name);
}
}

```

Hierarchy Level

[edit access]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure a subscriber access profile that includes 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. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

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 Access Profile on the LNS

Configuring an L2TP LNS with Inline Service Interfaces

Configuring PPP Properties for a Client-Specific Profile

Configuring Service Accounting with JSRC

[Configuring Service Accounting in Local Flat Files | 163](#)

AAA Service Framework Overview

Enabling Direct PCC Rule Activation by a PCRF for Subscriber Management

protocols

Syntax

```

protocols {
  bgp {
    ... bgp-configuration ...
  }
  isis {
    ... isis-configuration ...
  }
  ldp {
    ... ldp-configuration ...
  }
  mpls {
    ... mpls-configuration ...
  }
  msdp {
    ... msdp-configuration ...
  }
  mstp {
    ... mstp-configuration ...
  }
  ospf {
    domain-id domain-id;
    domain-vpn-tag number;
    route-type-community (iana | vendor);
    traffic-engineering {
      <advertise-unnumbered-interfaces>;
      <credibility-protocol-preference>;
      ignore-lsp-metrics;
      multicast-rpf-routes;
      no-topology;
      shortcuts {
        lsp-metric-into-summary;
      }
    }
    ... ospf-configuration ...
  }
  ospf3 {
    domain-id domain-id;
    domain-vpn-tag number;
    route-type-community (iana | vendor);
    traffic-engineering {
      <advertise-unnumbered-interfaces>;

```

```

    <credibility-protocol-preference>;
    ignore-lsp-metrics;
    multicast-rpf-routes;
    no-topology;
    shortcuts {
        lsp-metric-into-summary;
    }
}
... ospf3-configuration ...
}
pim {
    ... pim-configuration ...
}
rip {
    ... rip-configuration ...
}
ripng {
    ... ripng-configuration ...
}
rstp {
    rstp-configuration;
}
rsvp{
    ... rsvp-configuration ...
}
vstp {
    vstp configuration;
}
vpls {
    vpls configuration;
}
}

```

Hierarchy Level

```

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

```

Release Information

Statement introduced before Junos OS Release 7.4.

Support for RIPng introduced in Junos OS Release 9.0.

Statement introduced in Junos OS Release 11.1 for EX Series switches.

Statement introduced in Junos OS Release 11.3 for the QFX Series.

mpls and **rsvp** options added in Junos OS Release 15.1.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Specify the protocol for a routing instance. You can configure multiple instances of many protocol types.

Not all protocols are supported on the switches. See the switch CLI.

Options

bgp—Specify BGP as the protocol for a routing instance.

isis—Specify IS-IS as the protocol for a routing instance.

ldp—Specify LDP as the protocol for a routing instance or for a virtual router instance.

l2vpn—Specify Layer 2 VPN as the protocol for a routing instance.

mpls—Specify MPLS as the protocol for a routing instance.

msdp—Specify the Multicast Source Discovery Protocol (MSDP) for a routing instance.

mstp—Specify the Multiple Spanning Tree Protocol (MSTP) for a virtual switch routing instance.

ospf—Specify OSPF as the protocol for a routing instance.

ospf3—Specify OSPF version 3 (OSPFv3) as the protocol for a routing instance.

NOTE: OSPFv3 supports the **no-forwarding**, **virtual-router**, and **vrf** routing instance types only.

pim—Specify the Protocol Independent Multicast (PIM) protocol for a routing instance.

rip—Specify RIP as the protocol for a routing instance.

ripng—Specify RIP next generation (RIPng) as the protocol for a routing instance.

rstp—Specify the Rapid Spanning Tree Protocol (RSTP) for a virtual switch routing instance.

rsvp—Specify the RSVP for a routing instance.

vstp—Specify the VLAN Spanning Tree Protocol (VSTP) for a virtual switch routing instance.

vpls—Specify VPLS as the protocol for a routing instance.

Required Privilege Level

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| *Example: Configuring Multiple Routing Instances of OSPF*

proxy-arp

Syntax

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

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.6 for EX Series switches.

restricted added in Junos OS Release 10.0 for EX Series switches.

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description

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

NOTE: You must configure the IP address and the **inet** family for the interface when you enable proxy ARP.

Default

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

Options

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

Default: **unrestricted**

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Restricted and Unrestricted Proxy ARP

Configuring Proxy ARP on Switches

Example: Configuring Proxy ARP on an EX Series Switch

Configuring Gratuitous ARP

proxy-arp (Dynamic Profiles)

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

push (Dynamic VLANs)

Syntax

```
push;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number input-vlan-map]
```

Release Information

Statement introduced in Junos OS Release 10.4.

Description

For dynamic VLAN interfaces, specify the VLAN rewrite operation to add a new VLAN tag to the top of the VLAN stack. An outer VLAN tag is pushed in front of the existing VLAN tag. If you include the **push** statement in the configuration, you must also include the *pop* statement at the `[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number output-vlan-map]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Stacking and Rewriting VLAN Tags for the Layer 2 Wholesale Solution](#) | 94

push-backup-to-master (Accounting Options)

Syntax

```
push-backup-to-master;
```

Hierarchy Level

[edit accounting-options *file filename*]

Release Information

Statement introduced in Junos OS Release 16.1R4.

Description

Configure the router to save the accounting files from the new backup Routing Engine to the new master Routing Engine when a change in mastership occurs. The files are saved to the **/var/log/pfedBackup** directory on the router. The master Routing Engine includes these accounting files with its own current accounting files when it transfers the files from the backup directory to the archive site at the next transfer interval. Use this statement when the new backup Routing Engine is not able to connect to the archive site; for example, when site is not connected by means of an out-of-band interface or the path to the site is routed through a line card.

NOTE: The backup Routing Engine's files on the master Routing Engine are sent at each interval even though the files remain the same. If this is more activity than you want, consider using the **backup-on-failure master-and-slave** statement instead.

Required Privilege Level

system—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Accounting-Data Log Files

[Configuring Flat-File Accounting for Layer 2 Wholesale | 153](#)

[Configuring Flat-File Accounting for Extensible Subscriber Services Management | 158](#)

[Flat-File Accounting Overview | 149](#)

radius (Access Profile)

Syntax

```
radius {
  accounting-server [ ip-address ];
  attributes {
    exclude
      attribute-name packet-type;
    standard-attribute number {
      packet-type [ access-request | accounting-off | accounting-on | accounting-start | accounting-stop ];
    }
    vendor-id id-number {
      vendor-attribute vsa-number {
        packet-type [ access-request | accounting-off | accounting-on | accounting-start | accounting-stop ];
      }
    }
  }
  ignore {
    dynamic-iflset-name;
    framed-ip-netmask;
    idle-timeout;
    input-filter;
    logical-system-routing-instance;
    output-filter;
    session-timeout;
    standard-attribute number;
    vendor-id id-number {
      vendor-attribute vsa-number;
    }
  }
}
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;
  }
  chap-challenge-in-request-authenticator;
  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-channel;
    exclude-sub-interface;
}
ip-address-change-notify message;
juniper-access-line-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;
    interface-text-description;
    nas-identifier;
    order {
        agent-circuit-id;
        agent-remote-id;
        interface-description;
        interface-text-description;
        nas-identifier;
        postpend-vlan-tags;
    }
    postpend-vlan-tags;
}

```

```

nas-port-type {
    ethernet {
        port-type;
    }
}
override {
    calling-station-id remote-circuit-id;
    nas-ip-address tunnel-client-gateway-address;
    nas-port tunnel-client-nas-port;
    nas-port-type tunnel-client-nas-port-type;
}
remote-circuit-id-delimiter;
remote-circuit-id-fallback;
remote-circuit-id-format {
    agent-circuit-id;
    agent-remote-id;
}
revert-interval interval;
service-activation {
    dynamic-profile (optional-at-login | required-at-login);
    extensible-service (optional-at-login | required-at-login);
}
vlan-nas-port-stacked-format;
}
preauthentication-server ip-address;
}

```

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.

Options

accounting-server—(MX Series only) Specify a list of the RADIUS accounting servers used for accounting for DHCP, L2TP, and PPP clients.

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

authentication-server—(SRX Series only) 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.

Values: *ip-address*—IPv4 address.

preauthentication-server—(MX Series only) Starting in Junos OS Release 13.3, specify the RADIUS preauthentication server, which is used for the LLID service.

NOTE: You cannot configure this statement if the Calling-Station-ID attribute is excluded from RADIUS Access-Request messages by the **exclude** statement.

Values: *ip-address*—IPv4 address.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

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](#)

[RADIUS Logical Line Identifier \(LLID\) Overview](#)

[Configuring Logical Line Identification \(LLID\) Preauthentication](#)

radius-server

Syntax

```
radius-server server-address {
    accounting-port port-number;
    accounting-retry number;
    accounting-timeout seconds;
    dynamic-request-port port-number;
    max-outstanding-requests value;
    port port-number;
    preauthentication-port port-number;
    preauthentication-secret password;
    retry attempts;
    routing-instance routing-instance-name;
    secret password;
    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.

max-outstanding-requests introduced in Junos OS Release 11.4.

accounting-retry and **accounting-timeout** introduced in Junos OS Release 14.1.

dynamic-request-port option added in Junos OS Release 14.2R1 for MX Series routers.

preauthentication-port and **preauthentication-secret** options added in Junos OS Release 15.1 for MX Series routers.

accounting-port introduced in Junos OS Release 13.2X50-D10 for EX Series switches with support for Enhanced Layer 2 software (ELS). It was introduced in Junos OS without ELS in the following releases: Junos OS Releases 12.3R10, 14.1X53-D25, and 15.1R4 for EX Series switches.

Support for IPv6 **server-address** introduced in Junos OS Release 16.1.

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—IPv4 or IPv6 address of the RADIUS server.

accounting-port—(EX Series, M Series, MX Series, PTX Series, T Series only) Configure the port number on which to contact the RADIUS accounting server. This statement was introduced in Junos OS Release 13.2X50-D10 for EX Series switches with support for Enhanced Layer 2 software (ELS). It was introduced in Junos OS without ELS in the following releases: Junos OS Releases 12.3R10, 14.1X53-D25, and 15.1R4 for EX Series switches.

NOTE: Specifying the accounting port is optional, and port 1813 is the default. However, we recommend that you configure it in order to avoid confusion, as some RADIUS servers might refer to an older default.

Values: *port-number*—Port number on which to contact the RADIUS accounting server. Most RADIUS servers use port 1813, as specified in RFC 2866.

Default: 1813

accounting-retry—(MX Series, T Series only) Starting in Junos OS Release 14.1, configure the number of times the device retransmits RADIUS accounting messages when no response is received from the server. When you do not configure this statement, the number of retry attempts is determined by the **retry** statement.

NOTE: To successfully set a retry limit for the accounting servers different from the authentication servers, you must configure both the **accounting-retry** and **accounting-timeout** statements. If you configure only one of these statements, then the value you configure is ignored in favor of the values configured with the **retry** and **timeout** statements.

NOTE: The maximum retry duration (the number of retries times the length of the timeout) cannot exceed 2700 seconds. An error message is displayed if you configure a longer duration.

Values: *number*—Number of retry attempts.

Range: 0 through 100

Default: 0 (disabled)

accounting-timeout—(MX Series, T Series only) Starting in Junos OS Release 14.1, configure how long the local device waits to receive a response from a RADIUS accounting server before retransmitting the message. When you do not configure this statement, the length of the timeout is determined by the **timeout** statement.

NOTE: To successfully set a timeout value for the accounting servers different from the authentication servers, you must configure both the **accounting-retry** and **accounting-timeout** statements. If you configure only one of these statements, then the value you configure is ignored in favor of the values configured with the **retry** and **timeout** statements.

NOTE: The maximum retry duration (the number of retries times the length of the timeout) cannot exceed 2700 seconds. An error message is displayed if you configure a longer duration.

Values: *seconds*—Duration of timeout period.

Range: 0 through 1000 seconds

Default: 0 (disabled)

dynamic-request-port—(MX Series only) Starting in Junos OS Release 14.2R1, specify the port that the router monitors for dynamic (CoA) requests from the specified RADIUS servers. You can configure a port globally or for a specific access profile.

You must either use the default port for all RADIUS servers or configure the same nondefault port for all RADIUS servers. This rule applies at both the global access and access profile levels.

NOTE: Any other configuration results in a commit check failure. Multiple port numbers—that is, different port numbers for different servers—are not supported.

Values: *port-number*—Number of the monitored port.

Default: 3799 (as specified in RFC 5176)

max-outstanding-requests—(MX Series only) Starting in Junos OS Release 11.4, configure the maximum number of outstanding requests for this RADIUS server. An increase in this value is immediate while a decrease is more gradual if the current number of outstanding requests exceeds the new value.

Values: *requests*—Maximum number of outstanding requests for this RADIUS server.

Range: 0 through 2000 outstanding requests per server

Default: 1000 outstanding requests per server

port—(EX Series, M Series, MX Series, SRX Series, T Series only) Configure the port number on which to contact the RADIUS server.

Values: *port-number*—Port number on which to contact the RADIUS server.

Default: 1812 (as specified in RFC 2865)

preauthentication-port—(MX Series only) Starting in Junos OS Release 15.1 for MX Series routers, configure the port number on which to contact the RADIUS server for logical line identification (LLID) preauthentication requests. If you do not configure a separate UDP port for preauthentication purposes, the same UDP port that you configure for authentication messages by including the **port port-number** statement is used.

Values: *port-number*—Port number used for preauthentication requests to contact the RADIUS server.

preauthentication-secret—(MX Series only) Starting in Junos OS Release 15.1 for MX Series routers, configure the password to use with the RADIUS server for LLID preauthentication requests. If you do not configure a separate UDP password for preauthentication purposes, the same password that you configure for authentication messages by including the **secret password** statement is used. The secret password used by the local router must match that used by the server.

Values: *password*—Password to use. To include spaces enclose the character string in quotation marks.

retry—(EX Series, M Series, MX Series, PTX Series, T Series only) Specify the number of times that the device is allowed to attempt to contact a RADIUS authentication or accounting server. You can override the retry limit for accounting servers with the **accounting-retry** statement.

NOTE: To successfully set a retry limit for the accounting servers different from the authentication servers, you must configure both the **accounting-retry** and **accounting-timeout** statements. If you configure only one of these statements, then the value you configure is ignored in favor of the values configured with the **retry** and **timeout** statements.

NOTE: The maximum retry duration (the number of retries times the length of the timeout) cannot exceed 2700 seconds. An error message is displayed if you configure a longer duration.

Values: *attempts*—Number of times that the router is allowed to attempt to contact a RADIUS server.

Range: 1 through 100

Default: 3

routing-instance—(SRX Series, vSRX only) Configure the routing instance used to send RADIUS packets to the RADIUS server.

Values: *routing-instance-name*—Routing instance name.

source-address—(SRX Series, vSRX only) Configure a source address for each configured RADIUS server. Each RADIUS request sent to a RADIUS server uses the specified source address. Support for IPv6 **source-address** was introduced in Junos OS Release 16.1.

Values: *source-address*—Valid IPv4 or IPv6 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.

timeout—(SRX Series, vSRX only) Configure the amount of time that the local device waits to receive a response from RADIUS authentication and accounting servers. You can override the timeout value for accounting servers with the **accounting-timeout** statement.

NOTE: To successfully set a timeout value for the accounting servers different from the authentication servers, you must configure both the **accounting-retry** and **accounting-timeout** statements. If you configure only one of these statements, then the value you configure is ignored in favor of the values configured with the **retry** and **timeout** statements.

NOTE: The maximum retry duration (the number of retries times the length of the timeout) cannot exceed 2700 seconds. An error message is displayed if you configure a longer duration.

Values: *seconds*—Amount of time to wait.

Range: 1 through 1000 seconds

Default: 3 seconds

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Authentication and Accounting Parameters for Subscriber Access

Configuring the PPP Authentication Protocol

Configuring RADIUS Authentication for L2TP

Configuring RADIUS System Accounting

Configuring RADIUS-Initiated Dynamic Request Support

RADIUS Logical Line Identifier (LLID) Overview

RADIUS Attributes for LLID Preauthentication Requests

[show network-access aaa statistics](#) | 804

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

Configuring a Named Address Range for Dynamic Address Assignment

Address-Assignment Pools Overview

Address-Assignment Pool Configuration Overview

ranges (Dynamic VLAN)

Syntax

```
ranges (any | low-tag)-(any | high-tag);
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges dynamic-profile profile-name]
```

Release Information

Statement introduced in Junos OS Release 9.5.

Description

Configure VLAN ranges for dynamic, auto-sensed VLANs.

Options

any—The entire VLAN range.

low-tag—The lower limit of the VLAN range.

high-tag—The upper limit of the VLAN range.

Range: 1 through 4094

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| [Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs](#)

remote-id (VLAN Authentication Username)

Syntax

```
remote-id;
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges authentication username-include]
```

Release Information

Statement introduced in Junos OS Release 16.1R4.

Description

Include the agent remote identifier (ARI) in the username sent to RADIUS for authentication of the dynamic VLAN. The ARI is conveyed by the Access-Loop-Remote-ID TLV in an out-of-band ANCP Port Up message.

NOTE: This statement is not supported for stacked VLANs.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring a Username for Authentication of Out-of-Band Triggered Dynamic VLANs | 139](#)

[Configuring VLAN Interface Username Information for AAA Authentication](#)

[Layer 2 Wholesale with ANCP-Triggered VLANs Overview | 116](#)

route-distinguisher

Syntax

```
route-distinguisher (as-number:id | ip-address:id);
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols l2vpn mesh-group
  mesh-group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls mesh-group
  mesh-group-name],
[edit routing-instances routing-instance-name],
[edit routing-instances routing-instance-name protocols l2vpn mesh-group mesh-group-name],
[edit routing-instances routing-instance-name protocols vpls mesh-group mesh-group-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 11.1 for EX Series switches.

Support at **[edit routing-instances *routing-instance-name* protocols vpls mesh-group *mesh-group-name*]**

hierarchy level introduced in Junos OS Release 11.2.

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

Support at **[edit routing-instances *routing-instance-name* protocols l2vpn mesh-group *mesh-group-name*]**

hierarchy level introduced in Junos OS Release 13.2.

Statement introduced in Junos OS Release 14.1X53-D30 for QFX Series switches.

Statement introduced in cRPD Release 19.4R1.

Description

Specify an identifier attached to a route, enabling you to distinguish to which VPN or virtual private LAN service (VPLS) the route belongs. Each routing instance must have a unique route distinguisher (RD) associated with it. The RD is used to place bounds around a VPN so that the same IP address prefixes can be used in different VPNs without having them overlap. If the instance type is **vrf**, the **route-distinguisher** statement is required.

For Layer 2 VPNs and VPLS, if you configure the **l2vpn-use-bgp-rules** statement, you must configure a unique RD for each PE router participating in the routing instance.

For other types of VPNs, we recommend that you use a unique RD for each provider edge (PE) router participating in specific routing instance. Although you can use the same RD on all PE routers for the same VPN routing instance, if you use a unique RD, you can determine the customer edge (CE) router from which a route originated within the VPN.

For Layer 2 VPNs and VPLSs, if you configure mesh groups, the RD in each mesh group must be unique.



CAUTION: We strongly recommend that if you change an RD that has already been configured, or change the routing-instance type from **virtual-router** to **vrf**, make the change during a maintenance window, as follows:

1. Deactivate the routing instance.
2. Change the RD.
3. Activate the routing instance.

This is not required if you are configuring the RD for the first time.

Options

as-number:number—as-number is an assigned AS number, and **number** is any 2-byte or 4-byte value. The AS number can be from 1 through 4,294,967,295. If the AS number is a 2-byte value, the administrative number is a 4-byte value. If the AS number is 4-byte value, the administrative number is a 2-byte value. An RD consisting of a 4-byte AS number and a 2-byte administrative number is defined as a type 2 RD in RFC 4364 *BGP/MPLS IP VPNs*.

NOTE: In Junos OS Release 9.1 and later, the numeric range for AS numbers is extended to provide BGP support for 4-byte AS numbers, as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*. All releases of Junos OS support 2-byte AS numbers. To configure an RD that includes a 4-byte AS number, append the letter “L” to the end of the AS number. For example, an RD with the 4-byte AS number 7,765,000 and an administrative number of 1,000 is represented as **77765000L:1000**.

In Junos OS Release 9.2 and later, you can also configure a 4-byte AS number using the AS dot notation format of two integer values joined by a period: *<16-bit high-order value in decimal>.<16-bit low-order value in decimal>*. For example, the 4-byte AS number of 65,546 in the plain-number format is represented as 1.10 in AS dot notation format.

ip-address:id—IP address (**ip-address** is a 4-byte value) within your assigned prefix range and a 2-byte value for the **id**. The IP address can be any globally unique unicast address.

Range: 0 through 4,294,967,295 ($2^{32} - 1$). If the router you are configuring is a BGP peer of a router that does not support 4-byte AS numbers, you need to configure a local AS number. For more information, see *Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview*.

NOTE: For Ethernet VPN (EVPN), an RD that includes zero as the **id** value is reserved for the default EVPN routing instance by default. Because the same RD cannot be assigned for two routing instances, using a **ip-address:id** RD for another routing instance (default-switch), where the **id** value is zero, throws a commit error.

Required Privilege Level

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

<i>Example: Configuring BGP Route Target Filtering for VPNs</i>
<i>Example: Configuring FEC 129 BGP Autodiscovery for VPWS</i>
<i>Configuring EVPN Routing Instances</i>
<i>Configuring Routing Instances on PE Routers in VPNs</i>
<i>Configuring an MPLS-Based Layer 2 VPN (CLI Procedure)</i>
<i>Configuring an MPLS-Based Layer 3 VPN (CLI Procedure)</i>
<i>path-selection</i>

routing-instances (Dynamic Profiles)

Syntax

```

routing-instances routing-instance-name {
  interface interface-name;
  multicast-snooping-options {
  }
  routing-options {
    access {
      route prefix {
        metric route-cost;
        next-hop next-hop;
        preference route-distance;
        tag route-tag;
        tag2 route-tag2;
      }
    }
  }
  access-internal {
    route subscriber-ip-address {
      qualified-next-hop underlying-interface {
        mac-address address;
      }
    }
  }
  multicast {
    interface interface-name {
      no-qos-adjust;
    }
  }
  rib routing-table-name {
    access {
      route prefix {
        metric route-cost;
        next-hop next-hop;
        preference route-distance;
        tag route-tag;
        tag2 route-tag2;
      }
    }
  }
  access-internal {
    route subscriber-ip-address {
      qualified-next-hop underlying-interface {
        mac-address address;
      }
    }
  }
}

```

```

    }
  }
}
}
}

```

Hierarchy Level

```

[edit dynamic-profiles]
[edit logical-systems logical-system-name ]

```

Release Information

Statement introduced in Junos OS Release 9.6.

Support at the **logical-systems** hierarchy level was introduced in Junos OS Release 14.2.

Description

Dynamically configure an additional routing entity for a router in a dynamic client profile or a dynamic service profile.

Options

routing-instance-name—The routing instance variable (*\$junos-routing-instance*). The routing instance variable is dynamically replaced with the routing instance the accessing client uses when connecting to the router.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| [Configuring a Dynamic Profile for use by a Retailer in the DHCPv4 Solution](#) | 22

routing-instances (Multiple Routing Entities)

Syntax

```
routing-instances routing-instance-name { ... }
```

Hierarchy Level

```
[edit],  
[edit logical-systems logical-system-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

remote-vtep-v6-list statement introduced in Junos OS Release 17.3 for MX Series routers with MPC and MIC interfaces.

Description

Configure an additional routing entity for a router. You can create multiple instances of BGP, IS-IS, OSPF, OSPFv3, and RIP for a router. You can also create multiple routing instances for separating routing tables, routing policies, and interfaces for individual wholesale subscribers (retailers) in a Layer 3 wholesale network.

Each routing instance consist of the following:

- A set of routing tables
- A set of interfaces that belong to these routing tables
- A set of routing option configurations

Each routing instance has a unique name and a corresponding IP unicast table. For example, if you configure a routing instance with the name **my-instance**, its corresponding IP unicast table is my-instance.inet.0. All routes for **my-instance** are installed into my-instance.inet.0.

Routes are installed into the default routing instance inet.0 by default, unless a routing instance is specified.

In Junos OS Release 9.0 and later, you can no longer specify a routing-instance name of *master*, *default*, or *bgp* or include special characters within the name of a routing instance.

In Junos OS Release 9.6 and later, you can include a slash (/) in a routing-instance name only if a logical system is not configured. That is, you cannot include the slash character in a routing-instance name if a logical system other than the default is explicitly configured. Routing-instance names, further, are restricted from having the form `__.*__` (beginning and ending with underscores). The colon : character cannot be used when multitopology routing (MTR) is enabled.

Default

Routing instances are disabled for the router.

Options

routing-instance-name—Name of the routing instance. This must be a non-reserved string of not more than 128 characters.

remote-vtep-list—Configure static remote VXLAN tunnel endpoints.

remote-vtep-v6-list—Configure static IPv6 remote VXLAN tunnel endpoints.

Required Privilege Level

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

<i>Example: Configuring Interprovider Layer 3 VPN Option A</i>
<i>Example: Configuring Interprovider Layer 3 VPN Option B</i>
<i>Example: Configuring Interprovider Layer 3 VPN Option C</i>

schema-version (Flat-File Accounting Options)

Syntax

```
schema-version schema-name;
```

Hierarchy Level

```
[edit accounting-options flat-file-profile profile-name]
```

Release Information

Statement introduced in Junos OS Release 16.1R4.

Description

Specify the name of the XML schema that defines the contents and format of the accounting file, and appears in the accounting record header.

Options

schema-name—Name of the schema.

Required Privilege Level

system—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Flat-File Accounting for Layer 2 Wholesale | 153](#)

[Configuring Flat-File Accounting for Extensible Subscriber Services Management | 158](#)

[Flat-File Accounting Overview | 149](#)

secret (RADIUS)

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

RADIUS Authentication and Accounting Server Definition

Example: Configuring CHAP Authentication with RADIUS

Configuring RADIUS Authentication for L2TP

Configuring the RADIUS Disconnect Server for L2TP

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 PPPoE Dynamic Profile

Subscriber Interfaces and PPPoE Overview

server-group

Syntax

```
server-group {
  server-group-name {
    server-ip-address;
  }
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name forwarding-options dhcp-relay],
[edit logical-systems logical-system-name forwarding-options dhcp-relay dhcpv6],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay
  dhcpv6],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay dhcpv6]
```

Release Information

Statement introduced in Junos OS Release 8.3.

Support at the [\[edit ... dhcpv6\]](#) hierarchy levels introduced in Junos OS Release 11.4.

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

Description

Specify the name of a group of DHCP server addresses for use by the extended DHCP relay agent. Apply the group with the [active-server-group](#) statement globally for all interfaces or for a named group of interfaces configured with the [group](#) statement. This mechanism enables you to apply different DHCP relay configurations for different groups of servers, with a common configuration for the servers within a server group.

Options

server-group-name—Name of the group of DHCP or DHCPv6 server addresses.

server-ip-address—IP address of the DHCP server belonging to this named server group. Use IPv6 addresses when configuring DHCPv6 support. Starting in Junos OS Release 18.4R1, you can configure up to 32 server IP addresses per group for DHCPv4 servers. In earlier releases, you can configure only up to 5 server IP addresses for DHCPv4 servers. For DHCPv6 servers, you can configure only up to 32 addresses in all releases. The configuration fails commit check if you configure more than the maximum number of server addresses.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Extended DHCP Relay Agent Overview

Configuring Active Server Groups to Apply a Common DHCP Relay Agent Configuration to Named Server Groups

site (VPLS Multihoming for FEC 128)

Syntax

```
site site-name {
  mac-pinning;
  active-interface (any | primary interface-name);
  best-site;
  interface interface-name {
    interface-mac-limit limit;
  }
  mesh-group mesh-group-name;
  multi-homing;
  site-identifier identifier;
  site-preference preference-value;
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the site name and site identifier for a site. Allows you to configure a remote site ID for remote sites.

Options

site-name—Name of the site.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

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

site-identifier (VPLS)

Syntax

```
site-identifier identifier;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls site site-name],  
[edit routing-instances routing-instance-name protocols vpls site site-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the numerical identifier for the local VPLS site.

Options

identifier—Specify the numerical identifier for the local VPLS site. The identifier must be an unsigned 16-bit number greater than zero.

Required Privilege Level

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| *Configuring the VPLS Site Name and Site Identifier*

site-range

Syntax

```
site-range number;
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify an upper limit on the maximum site identifier that can be accepted to allow a pseudowire to be brought up. Pseudowires cannot be established to sites with site identifiers greater than the configured site range. If you issue the **show vpls connections** command, such sites are displayed as OR (out of range).

Options

number—Maximum number of site identifiers. We recommend using the default value.

Range: 1 through 65,534

Default: 65,534

Required Privilege Level

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| *Configuring the Site Range*

stacked-vlan-ranges

Syntax

```
stacked-vlan-ranges {
  access-profile profile-name;
  authentication {
    packet-types [packet-types];
    password password-string;
    username-include {
      circuit-type;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-name;
      mac-address;
      option-18
      option-37
      option-82;
      radius-realm radius-realm-string;
      user-prefix user-prefix-string;
      vlan-tags;
    }
  }
  dynamic-profile profile-name {
    accept (any | dhcp-v4 | inet);
    access-profile vlan-dynamic-profile-name;
    ranges (any | low-tag-high-tag), (any | low-tag-high-tag);
  }
  override;
}
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure]
```

Release Information

Statement introduced in Junos OS Release 9.5.

Description

Configure multiple VLANs. Each VLAN is assigned a VLAN ID number from the range.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring an Interface to Use the Dynamic Profile Configured to Create Stacked VLANs

Configuring Interfaces to Support Both Single and Stacked VLANs

stacked-vlan-tagging**Syntax**

```
stacked-vlan-tagging;
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description

For Gigabit Ethernet IQ interfaces, Gigabit Ethernet, 10-Gigabit Ethernet LAN/WAN PIC, and 100-Gigabit Ethernet Type 5 PIC with CFP, enable stacked VLAN tagging for all logical interfaces on the physical interface.

For pseudowire subscriber interfaces, enable stacked VLAN tagging for logical interfaces on the pseudowire service.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview

traceoptions (DHCP)

Syntax

```
traceoptions {
  file filename <files number> <match regular-expression > <size maximum-file-size> <world-readable |
    no-world-readable>;
  flag flag;
  level (all | error | info | notice | verbose | warning);
  no-remote-trace;
}
```

Hierarchy Level

```
[edit system processes dhcp-service]
[edit security dynamic-address]
```

Release Information

Statement introduced in Junos OS Release 11.4.

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

Statement introduced in Junos OS Release 18.4R1.

Description

Define global tracing operations for extended DHCP local server and extended DHCP relay agent processes.

This statement replaces the deprecated **traceoptions** statements at the **[edit forwarding-options dhcp-relay]** and **[edit system services dhcp-local-server]** hierarchy levels.

NOTE: Traceoptions does not differentiate between a logical system and tenant system, and can be configured under the root logical system.

Options

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files *number*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements:

- **all**—Trace all events.
- **auth**—Trace authentication events.
- **database**—Trace database events.
- **fwd**—Trace firewall process events.
- **general**—Trace miscellaneous events.
- **ha**—Trace high availability-related events.
- **interface**—Trace interface operations.
- **io**—Trace I/O operations.
- **liveness-detection**—Trace liveness detection operations.
- **packet**—Trace packet and option decoding operations.
- **performance**—Trace performance measurement operations.
- **profile**—Trace profile operations.
- **rpdp**—Trace routing protocol process events.
- **rtsock**—Trace routing socket operations.
- **security-persistence**—Trace security persistence events.
- **session-db**—Trace session database events.
- **state**—Trace changes in state.
- **statistics**—Trace baseline statistics.
- **ui**—Trace user interface operations.

level—Level of tracing to perform; also known as severity level. The option you configure enables tracing of events at that level and all higher (more restrictive) levels. You can specify any of the following levels:

- **all**—Match messages of all levels.
- **error**—Match error messages.
- **info**—Match informational messages.
- **notice**—Match notice messages about conditions requiring special handling.
- **verbose**—Match verbose messages. This is the lowest (least restrictive) severity level; when you configure **verbose**, messages at all higher levels are traced. Therefore, the result is the same as when you configure **all**.
- **warning**—Match warning messages.

Default: error

match *regular-expression*—(Optional) Refine the output to include lines that contain the regular expression.

no-remote-trace—Disable remote tracing.

no-world-readable—(Optional) Disable unrestricted file access, allowing only the user **root** and users who have the Junos OS **maintenance** permission to access the trace files.

size *maximum-file-size*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (***maximum-file-size*k**), megabytes (***maximum-file-size*m**), or gigabytes (***maximum-file-size*g**). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Range: 10,240 through 1,073,741,824

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| *Tracing Extended DHCP Operations*

underlying-interface (demux0)

Syntax

```
underlying-interface underlying-interface-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0 interface-name unit unit logical-unit-number demux-options]
```

Release Information

Statement introduced in Junos OS Release 9.3.

Support for aggregated Ethernet introduced in Junos OS Release 9.4.

Description

Configure the underlying interface on which the demultiplexing (demux) interface is running.



CAUTION: Before you make any changes to the underlying interface for a demux0 interface, you must ensure that no subscribers are currently present on that underlying interface. If any subscribers are present, you must remove them before you make changes.

Options

underlying-interface-name—Either the specific name of the interface on which the DHCP discover packet arrives or one of the following interface variables:

- **\$junos-underlying-interface** when configuring dynamic IP demux interfaces.
- **\$junos-interface-ifd-name** when configuring dynamic VLAN demux interfaces.

The variable is used to specify the underlying interface when a new demux interface is dynamically created. The variable is dynamically replaced with the underlying interface that DHCP supplies when the subscriber logs in.

NOTE: Logical demux interfaces are currently supported on Gigabit Ethernet, Fast Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles

Configuring Dynamic Subscriber Interfaces Using VLAN Demux Interfaces in Dynamic Profiles

Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview

For information about static underlying interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*

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 PPPoE Dynamic Profile

Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview

unit

Syntax

```

unit logical-unit-number {
  accept-source-mac {
    mac-address mac-address {
      policer {
        input cos-policer-name;
        output cos-policer-name;
      }
    }
  }
  accounting-profile name;
  advisory-options {
    downstream-rate rate;
    upstream-rate rate;
  }
  allow-any-vci;
  atm-scheduler-map (map-name | default);
  auto-configure {
    agent-circuit-identifier {
      dynamic-profile profile-name;
    }
    line-identity {
      include {
        accept-no-ids;
        circuit-id;
        remote-id;
      }
      dynamic-profile profile-name;
    }
  }
  backup-options {
    interface interface-name;
  }
  bandwidth rate;
  cell-bundle-size cells;
  clear-dont-fragment-bit;
  compression {
    rtp {
      maximum-contexts number <force>;
      f-max-period number;
      queues [queue-numbers];
      port {

```

```
        minimum port-number;  
        maximum port-number;  
    }  
}  
}  
compression-device interface-name;  
copy-tos-to-outer-ip-header;
```

```

demux {
  inet {
    address-source address;
    auto-configure {
      address-ranges {
        authentication {
          password password-string;
          username-include {
            auth-server-realm realm-string;
            delimiter delimiter-character;
            domain-name domain-name;
            interface-name;
            source-address;
            user-prefix user-prefix-string;
          }
        }
      }
      dynamic-profile profile-name {
        network ip-address {
          range name {
            low lower-limit;
            high upper-limit;
          }
        }
      }
    }
  }
}

inet6 {
  address-source address;
  auto-configure {
    address-ranges {
      authentication {
        password password-string;
        username-include {
          auth-server-realm realm-string;
          delimiter delimiter-character;
          domain-name domain-name;
          interface-name;
          source-address;
          user-prefix user-prefix-string;
        }
      }
    }
    dynamic-profile profile-name {
      network ip-address {

```



```

dynamic-call-admission-control {
    activation-priority priority;
    bearer-bandwidth-limit kilobits-per-second;
}
encapsulation type;
epd-threshold cells plp1 cells;
family family-name {
    ... the family subhierarchy appears after the main [edit interfaces interface-name unit logical-unit-number] hierarchy
    ...
}
fragment-threshold bytes;
host-prefix-only;
inner-vlan-id-range start start-id end end-id;
input-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap |
    swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
}
interleave-fragments;
inverse-arp;
layer2-policer {
    input-policer policer-name;
    input-three-color policer-name;
    output-policer policer-name;
    output-three-color policer-name;
}
link-layer-overhead percent;
minimum-links number;
mrru bytes;
multicast-dlci dlci-identifier;
multicast-vci vpi-identifier.vci-identifier;
multilink-max-classes number;
multipoint;
oam-liveness {
    up-count cells;
    down-count cells;
}
oam-period (disable | seconds);

```

```

output-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap |
    swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
}
passive-monitor-mode;
peer-unit unit-number;
plp-to-clp;
point-to-point;
ppp-options {
    mru size;
    mtu (size | use-lower-layer);
    chap {
        access-profile name;
        default-chap-secret name;
        local-name name;
        passive;
    }
    compression {
        acfc;
        pfc;
    }
    dynamic-profile profile-name;
    ipcp-suggest-dns-option;
    lcp-restart-timer milliseconds;
    loopback-clear-timer seconds;
    ncp-restart-timer milliseconds;
    pap {
        access-profile name;
        default-pap-password password;
        local-name name;
        local-password password;
        passive;
    }
}
pppoe-options {
    access-concentrator name;
    auto-reconnect seconds;
    (client | server);
    service-name name;
    underlying-interface interface-name;
}

```

```

pppoe-underlying-options {
    access-concentrator name;
    direct-connect;
    dynamic-profile profile-name;
    max-sessions number;
}
proxy-arp;
service-domain (inside | outside);
shaping {
    (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate burst length);
    queue-length number;
}
short-sequence;
targeted-distribution;
transmit-weight number;
(traps | no-traps);
trunk-bandwidth rate;
trunk-id number;
tunnel {
    backup-destination address;
    destination address;
    key number;
    routing-instance {
        destination routing-instance-name;
    }
    source source-address;
    ttl number;
}
vci vpi-identifier.vci-identifier;
vci-range start start-vci end end-vci;
vpi vpi-identifier;
vlan-id number;
vlan-id-range number-number;
vlan-tags inner tpid.vlan-id outer tpid.vlan-id;

```

```

family family {
  accounting {
    destination-class-usage;
    source-class-usage {
      (input | output | input output);
    }
  }
  access-concentrator name;
  address address {
    ... the address subhierarchy appears after the main [edit interfaces interface-name unit logical-unit-number
      family family-name] hierarchy ...
  }
  bundle interface-name;
  core-facing;
  demux-destination {
    destination-prefix;
  }
  demux-source {
    source-prefix;
  }
  direct-connect;
  duplicate-protection;
  dynamic-profile profile-name;
  filter {
    group filter-group-number;
    input filter-name;
    input-list [filter-names];
    output filter-name;
    output-list [filter-names];
  }
  interface-mode (access | trunk);
  ipsec-sa sa-name;
  keep-address-and-control;
  mac-validate (loose | strict);
  max-sessions number;
  mtu bytes;
  multicast-only;
  no-redirects;
  policer {
    arp policer-template-name;
    input policer-template-name;
    output policer-template-name;
  }
  primary;

```

```

protocols [inet iso mpls];
proxy inet-address address;
receive-options-packets;
receive-ttl-exceeded;
remote (inet-address address | mac-address address);
rpf-check {
    fail-filter filter-name
    mode loose;
}
sampling {
    input;
    output;
}
service {
    input {
        post-service-filter filter-name;
        service-set service-set-name <service-filter filter-name>;
    }
    output {
        service-set service-set-name <service-filter filter-name>;
    }
}
service-name-table table-name
targeted-options {
    backup backup;
    group group;
    primary primary;
    weight ($junos-interface-target-weight | weight-value);
}
(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 dlcid-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.

Range increased for static pseudowire interfaces to 1,073,741,823 in Junos OS Release 18.3R1.

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, PPPoE, and pseudowire static interfaces. 0 through 16,385 for all other static interface types.

etree-ac-role (leaf | root)—To configure an interface as either leaf or root.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Logical Interface Properties

Junos OS Services Interfaces Library for Routing Devices

unit (Dynamic Demux Interface)

Syntax

```

unit logical-unit-number {
  demux-options {
    underlying-interface interface-name
  }
  family family {
    access-concentrator name;
    address address;
    demux-source {
      source-address;
    }
    direct-connect;
    duplicate-protection;
    dynamic-profile profile-name;
    filter {
      input filter-name;
      output filter-name;
    }
    mac-validate (loose | strict);
    max-sessions number;
    max-sessions-vsa-ignore;
    rpf-check {
      fail-filter filter-name;
      mode loose;
    }
    service-name-table table-name;
    short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max maximum-seconds>;
    unnumbered-address interface-name <preferred-source-address address>;
  }
  filter {
    input filter-name;
    output filter-name;
  }
}
vlan-id number;

```

Hierarchy Level

[edit [dynamic-profiles](#) profile-name [interfaces](#) demux0]

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Configure a dynamic logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options

logical-unit-number—Either the specific unit number of the interface or the unit number variable (**\$junos-interface-unit**). The variable is used to specify the unit of the interface when a new demux interface is dynamically created. The static unit number variable is dynamically replaced with the unit number that DHCP supplies when the subscriber logs in.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| *Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles*

unit (Dynamic Profiles Standard Interface)

Syntax

```

unit logical-unit-number {
    actual-transit-statistics;
    auto-configure {
        agent-circuit-identifier {
            dynamic-profile profile-name;
        }
        line-identity {
            include {
                accept-no-ids;
                circuit-id;
                remote-id;
            }
            dynamic-profile profile-name;
        }
    }
    dial-options {
        ipsec-interface-id name;
        l2tp-interface-id name;
        (shared | dedicated);
    }
    encapsulation (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-cisco-nlpid | atm-tcc-vc-mux | atm-mlppp-llc | atm-nlpid
    | atm-ppp-llc | atm-ppp-vc-mux | atm-snap | atm-tcc-snap | atm-vc-mux | ether-over-atm-llc |
    ether-vpls-over-atm-llc | ether-vpls-over-fr | ether-vpls-over-ppp | ethernet | frame-relay-ccc | frame-relay-ppp
    | frame-relay-tcc | frame-relay-ether-type | frame-relay-ether-type-tcc | multilink-frame-relay-end-to-end |
    multilink-ppp | ppp-over-ether | ppp-over-ether-over-atm-llc | vlan-bridge | vlan-ccc | vlan-vci-ccc | vlan-tcc |
    vlan-vpls);
    family family {
        address address;
        demux-destination,
        filter {
            adf {
                counter;
                input-precedence precedence;
                not-mandatory;
                output-precedence precedence;
                rule rule-value;
            }
            input filter-name {
                precedence precedence;
                shared-name filter-shared-name;
            }
        }
    }
}

```

```

    output filter-name {
        precedence precedence;
        shared-name filter-shared-name;
    }
}
max-sessions number;
max-sessions-vsa-ignore;
rpf-check {
    fail-filter filter-name;
    mode loose;
}
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 {
    shared-name filter-shared-name;
  }
  output filter-name {
    shared-name filter-shared-name;
  }
}
host-prefix-only;
keepalives {
  interval seconds;
}
ppp-options {
  aaa-options aaa-options-name;
  authentication [ authentication-protocols ];
  chap {
    challenge-length minimum minimum-length maximum maximum-length;
    local-name name;
  }
  ignore-magic-number-mismatch;
  initiate-ncp (dual-stack-passive | ipv6 | ip)
  ipcp-suggest-dns-option;
  mru size;
  mtu (size | use-lower-layer);
  on-demand-ip-address;
  pap;
  peer-ip-address-optional;
  local-authentication {
    password password;
    username-include {
      circuit-id;
      delimiter character;
      domain-name name;
      mac-address;
      remote-id;
    }
  }
}
service {
  pcef pcef-profile-name {
    activate rule-name | activate-all;
  }
}

```

```

targeted-options {
    backup backup;
    group group;
    primary primary;
    weight ($junos-interface-target-weight | weight-value);
}
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 predefined variables:

- **\$junos-underlying-interface-unit**—For static VLANs, the unit number variable. The static unit number variable is dynamically replaced with the client unit number when the client session begins. The client unit number is specified by the DHCP when it accesses the subscriber network.
- **\$junos-interface-unit**—The unit number variable on a dynamic underlying VLAN interface for which you want to enable the creation of dynamic VLAN subscriber interfaces based on the ACI.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Dynamic Underlying VLAN Interfaces to Use Agent Circuit Identifier Information](#)

[Configuring Static Underlying VLAN Interfaces to Use Agent Circuit Identifier Information](#)

[Agent Circuit Identifier-Based Dynamic VLANs Overview](#)

unnumbered-address (Dynamic PPPoE)

Syntax

```
unnumbered-address interface-name;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces pp0 unit "$junos-interface-unit" family inet]
```

Release Information

Statement introduced in Junos OS Release 10.1.

Description

For dynamic PPPoE 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.

Options

interface-name—Interface from which the local address is derived. The interface name must include a logical unit number and must have a configured address.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring a PPPoE Dynamic Profile](#)

[Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview](#)

unnumbered-address (Dynamic Profiles)

Syntax

```
unnumbered-address interface-name <preferred-source-address address>;
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 9.2.

Support for the **\$junos-preferred-source-address** and **\$junos-preferred-source-ipv6-address** predefined variables introduced in Junos OS Release 9.6.

Support for the **\$junos-loopback-interface** predefined variable introduced in Junos OS Release 9.6.

Description

For Ethernet interfaces, enable the local address to be derived from the specified interface. Configuring unnumbered Ethernet interfaces enables IP processing on the interface without assigning an explicit IP address to the interface. To configure unnumbered address dynamically, include the **\$junos-loopback-interface-address** predefined variable.

You can configure unnumbered address support on Ethernet interfaces for IPv4 and IPv6 address families.

Options

interface-name—Name of the interface from which the local address is derived. The specified interface must have a logical unit number, a configured IP address, and must not be an unnumbered interface. This value can be a specific interface name or the **\$junos-loopback-interface** predefined variable.

When defining the **unnumbered-address** statement using a static interface, keep the following in mind:

- If you choose to include the **routing-instance** statement at the **[edit dynamic-profiles]** hierarchy level, that statement must be configured with a dynamic value by using the **\$junos-routing-instance** predefined variable. In addition, whatever static unnumbered interface you specify must belong to that routing instance; otherwise, the profile instantiation fails.
- If you choose to not include the **routing-instance** statement at the **[edit dynamic-profiles]** hierarchy level, the unnumbered-address statement uses the default routing instance. The use of the default routing instance requires that the unnumbered interface be configured statically and that it reside in the default routing instance.

NOTE: When you specify a static logical interface for the unnumbered interface in a dynamic profile that includes the **\$junos-routing-instance** predefined variable, you must not configure a preferred source address, whether with the **\$junos-preferred-source-address** predefined variable, the **\$junos-preferred-source-ipv6-address** predefined variable, or the **preferred-source-address** statement. Configuring the preferred source address in this circumstance causes a commit failure.

When defining the **unnumbered-address** statement using the **\$junos-loopback-interface** predefined variable, keep the following in mind:

- To use the **\$junos-loopback-interface** predefined variable, the dynamic profile must also contain the **routing-instance** statement configured with the **\$junos-routing-instance** predefined variable at the [edit dynamic-profiles] hierarchy level.
- The applied loopback interface is based on the dynamically obtained routing instance of the subscriber.

address—(Optional) Secondary IP address of the donor interface. Configuring the preferred source address enables you to use an IP address other than the primary IP address on some of the unnumbered Ethernet interfaces in your network. This value can be a static IP address, the **\$junos-preferred-source-address** predefined variable for the inet family, or the **\$junos-preferred-source-ipv6-address** predefined variable for the inet6 family.

When defining the **preferred-source-address** value using a static IP address, keep the following in mind:

- The unnumbered interface must be statically configured.
- The IP address specified as the **preferred-source-address** must be configured in the specified unnumbered interface.

When defining the **preferred-source-address** value using the **\$junos-preferred-source-address** or the **\$junos-preferred-source-ipv6-address** predefined variables, keep the following in mind:

- You must configure the **unnumbered-address** statement using the **\$junos-loopback-interface** predefined variable.
- You must configure the **routing-instance** statement using the **\$junos-routing-instance** predefined variable at the [edit dynamic-profiles] hierarchy level.
- The preferred source address chosen is based on the dynamically applied loopback address which is in turn derived from the dynamically obtained routing instance of the subscriber. The configured loopback address with the closest network match to the user IP address is selected as the preferred source address.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Dynamic Profiles Overview](#)

unnumbered-address (Ethernet)

Syntax

```
unnumbered-address interface-name <preferred-source-address address>;
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 8.2.

preferred-source-address option introduced in Junos OS Release 9.0.

Description

For Ethernet interfaces, enable the local address to be derived from the specified interface. Configuring an unnumbered Ethernet interface enables IP processing on the interface without assigning an explicit IP address to the interface.

Options

interface-name—Name of the interface from which the local address is derived. The specified interface must have a logical unit number and a configured IP address, and must not be an unnumbered interface.

The **preferred-source-address** statement is explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring an Unnumbered Interface](#)

[address](#) | 207

[Junos System Basics Configuration Guide](#)

username-include (Interfaces)

Syntax

```
username-include {
  circuit-id;
  circuit-type;
  delimiter delimiter-character;
  domain-name domain-name-string;
  interface-name;
  mac-address;
  option-18;
  option-37;
  option-82 <circuit-id> <remote-id>;
  radius-realm radius-realm-string;
  remote-id;
  user-prefix user-prefix-string;
  vlan-tags;
}
```

Hierarchy Level

```
[edit interfaces interface-name auto-configure vlan-ranges authentication],
[edit interfaces interface-name auto-configure stacked-vlan-ranges authentication]
```

Release Information

Statement introduced in Junos OS Release 10.0.

vlan-tags option added in Junos OS Release 18.3R1 on MX Series routers.

Description

Configure the username that the router passes to the external AAA server. You must include at least one of the optional statements for the username to be valid. If you do not configure a username, the router accesses the local authentication service only and does not use external authentication services, such as RADIUS.

The username takes the format ***user-prefix mac-address circuit-type circuit-id remote-id option-82 interface-name domain-name radius-realm***. By default, each component is separated by a period (.), but you can specify a different delimiter with the **delimiter** statement.

Options

vlan-tags—Include the subscriber session VLAN tags in the username for interactions with an external authority. Both single-tagged and double-tagged VLANs are supported: The tags are added in the

format ***outer-vlan-tag-inner-vlan-tag***. The outer tag is always included; the inner tag is included for double-tagged VLANs.

You can use this option instead of the **interface-name** option when the outer VLAN tag is unique across the system and you do not need the underlying physical interface name to be part of the format.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring VLAN Interface Username Information for AAA Authentication

Using DHCP Option 82 Suboptions in Authentication Usernames for Autosense VLANs

Using DHCP Option 18 and Option 37 in Authentication Usernames for DHCPv6 Autosense VLANs

[Configuring a Username for Authentication of Out-of-Band Triggered Dynamic VLANs | 139](#)

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

| *Specifying Authentication Support*

vlan-id (Dynamic VLANs)

Syntax

```
vlan-id number;
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number input-vlan-map],  
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number output-vlan-map]
```

Release Information

Statement introduced in Junos OS Release 10.4.

Description

For dynamic VLAN interfaces, specify the line VLAN identifiers to be rewritten at the input or output interface.

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

Options

number—A valid VLAN identifier. When used for input VLAN maps, you can specify the **\$junos-vlan-map-id** predefined variable to dynamically obtain the VLAN identifier.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Rewriting the VLAN Tag on Tagged Frames

Binding VLAN IDs to Logical Interfaces

vlan-model

Syntax

```
vlan-model one-to-one;
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 11.2.

Description

Define the network VLAN model.

Options

one-to-one—Specify that any received, dual-tagged VLAN packet triggers the provisioning process in a Layer 2 Wholesale network. Using this option, the router learns VLAN tags for each individual client. The router learns both the outer tag and inner tag of the incoming packets, when the **instance-role** statement is defined as **access**, or the outer VLAN tag only, when the **instance-role** statement is defined as **nni**.

Required Privilege Level

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Separate Access Routing Instances for Layer 2 Wholesale Service Retailers | 102](#)

[Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers | 105](#)

vlan-ranges

Syntax

```
vlan-ranges {
  access-profile profile-name;
  authentication {
    packet-types [packet-types];
    password password-string;
    username-include {
      circuit-type;
      circuit-id;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-name;
      mac-address;
      option-18;
      option-37;
      option-82 <circuit-id> <remote-id>;
      radius-realm radius-realm-string;
      remote-id;
      user-prefix user-prefix-string;
      vlan-tags;
    }
  }
  dynamic-profile profile-name {
    accept (any | dhcp-v4 | inet);
    accept-out-of-band protocol;
    access-profile vlan-dynamic-profile-name;
    ranges (any | low-tag)–(any | high-tag);
  }
  override;
}
```

Hierarchy Level

[edit interfaces *interface-name* **auto-configure**]

Release Information

Statement introduced in Junos OS Release 9.5.

Description

Configure multiple VLANs. Each VLAN is assigned a VLAN ID number from the range.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring an Interface to Use the Dynamic Profile Configured to Create Single-Tag VLANs

Configuring Interfaces to Support Both Single and Stacked VLANs

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 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.

NOTE: On the network-to-network (NNI) or egress interfaces of provider edge (PE) routers, you cannot configure the **inner-range tpid. vid1–vid2** option with the **vlan-tags** statement for ISP-facing interfaces.

Range: For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

outer [tpid].vlan-id—A TPID (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*

vlan-tags (Stacked VLAN Tags)

Syntax

```
vlan-tags inner tpid.vlan-id inner-list value inner-range vid1—vid2 outer tpid.vlan-id;
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

Description

Bind TPIDs and 802.1Q VLAN tag IDs to a logical interface. TPID fields are used to identify the frame as an IEEE 802.1Q-tagged frame.

Options

inner *tpid.vlan-id*—A TPID and a valid VLAN identifier. TPID is a 16-bit field set to a value of 0x8100 in order to identify the frame as an IEEE 802.1Q-tagged frame.

Range: (most routers) For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames. For PTX Series, VLAN ID 0 is not supported.

inner-list *value*— List or a set of VLAN identifiers.

NOTE: This is supported on MX Series routers with Trio-based FPCs.

inner-range *tpid. vid1—vid2*—Specify a TPID and a range of VLAN IDs where vid1 is the start of the range and vid2 is the end of the range.

NOTE: On the network-to-network (NNI) or egress interfaces of provider edge (PE) routers, you cannot configure the **inner-range *tpid. vid1—vid2*** option with the **vlan-tags** statement for ISP-facing interfaces.

Range: For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

outer *tpid.vlan-id*—A TPID and a valid VLAN identifier.

Range: (most routers) For VLAN ID, 1 through 511 for normal interfaces, and 512 through 4094 for VLAN CCC interfaces. VLAN ID 0 is reserved for tagging the priority of frames. For PTX Series, VLAN ID 0 is not supported.

NOTE: Configuring **inner-range** with the entire vlan-id range consumes system resources and is not a best practice. The **inner-range** must be used only when a subset of VLAN IDs of inner tag (not the entire range) needs to be associated with a logical interface. If you specify the entire range (1 through 4094), it 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](#) | 459

vpls (Routing Instance)

Syntax

```
vpls {
  mac-pinning;
  active-interface {
    any;
    primary interface-name;
  }
  community COMM;
  connectivity-type (ce | irb);
  control-word;
  encapsulation-type ethernet;
  ignore-encapsulation-mismatch;
  ignore-mtu-mismatch;
  import-labeled-routes [ routing-instance-name ];
  interface interface-name;
  interface-mac-limit limit;
  label-block-size size;
  mac-flush [ explicit-mac-flush-message-options ];
  mac-table-aging-time time;
  mac-table-size size;
  mesh-group mesh-group-name {
    interface interface-name;
    l2vpn-id (as-number:id | ip-address:id);
    local-switching;
    mac-flush [ explicit-mac-flush-message-options ];
    neighbor address {...};
    peer-as all;
    pseudowire-status-tlv hot-standby-vc-on;
    route-distinguisher (as-number:id | ip-address:id);
    vpls-id number;
    vrf-export [ policy-names ];
    vrf-import [ policy-names ];
    vrf-target {
      community;
      import community-name;
      export community-name;
    }
  }
  mtu mtu;
  no-control-word;
  no-tunnel-services;
  site site-name {
```

```

active-interface interface-name {
    any;
    primary preference-value;
}
best-site;
interface interface-name {
    interface-mac-limit limit;
}
mesh-group mesh-group-name;
multi-homing;
site-identifier identifier;
site-preference preference-value {
    backup;
    primary;
}
}
site-range number;
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
tunnel-services {
    devices device-names;
    primary primary-device-name;
}
}

```

Hierarchy Level

```

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

```

Release Information

Statement introduced before Junos OS Release 7.4.

mac-flush option introduced in Junos OS Release 10.0.

hot-standby-vc-on, **import-labeled-routes** [*routing-instance-name*], and **interface** *interface* options introduced in Junos OS Release 15.1R2.

Description

Configure a virtual private LAN service (VPLS) routing instance.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| *Configuring VPLS Routing Instances*

vrf-export

Syntax

```
vrf-export [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls mesh-group  
  mesh-group-name]  
[edit routing-instances routing-instance-name]  
[edit routing-instances routing-instance-name protocols vpls mesh-group mesh-group-name]  
[edit switch-options]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 11.1 for EX Series switches.

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

Statement introduced in Junos OS Release 14.1X53-D30 for QFX Series switches.

Statement introduced in cRPD Release 19.4R1.

Description

Specify how routes are exported from the local PE router's VRF table (*routing-instance-name.inet.0*) to the remote PE router. If the value **vrf** is specified for the **instance-type** statement included in the routing instance configuration, this statement is required.

You can configure multiple export policies on the PE router or PE switch.

Default

If the instance-type is **vrf**, **vrf-export** is a required statement. The default action is to reject.

Options

policy-names— Names for the export policies.

Required Privilege Level

routing— To view this statement in the configuration.

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

RELATED DOCUMENTATION

Implementing EVPN-VXLAN for Data Centers

[instance-type](#) | 369

Configuring Policies for the VRF Table on PE Routers in VPNs

vrf-import

Syntax

```
vrf-import [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls mesh-group  
  mesh-group-name]  
[edit routing-instances routing-instance-name]  
[edit routing-instances routing-instance-name protocols vpls mesh-group mesh-group-name]  
[edit switch-options]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 11.1 for EX Series switches.

Statement introduced in Junos OS Release 14.1X53-D30 for QFX Series switches.

Statement introduced in cRPD Release 19.4R1.

Description

Specify how routes are imported into the virtual routing and forwarding (VRF) table (***routing-instance-name*.inet.0**) of the local provider edge (PE) router or switch from the remote PE router. If the value **vrf** is specified for the **instance-type** statement included in the routing instance configuration, this statement is required.

You can configure multiple import policies on the PE router or switch.

Default

If the instance type is **vrf**, **vrf-import** is a required statement. The default action is to accept.

Options

policy-names—Names for the import policies.

Required Privilege Level

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Implementing EVPN-VXLAN for Data Centers

[instance-type](#) | [369](#)

Configuring Policies for the VRF Table on PE Routers in VPNs

vrf-target

Syntax

```
vrf-target {
    community;
    auto
    import community-name;
    export community-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols l2vpn mesh-group
mesh-group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls mesh-group
mesh-group-name],
[edit routing-instances routing-instance-name protocols evpn vni-options],
[edit routing-instances routing-instance-name protocols l2vpn mesh-group mesh-group-name],
[edit routing-instances routing-instance-name protocols vpls mesh-group mesh-group-name],
[edit switch-options]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 11.1 for EX Series switches.

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

Statement introduced in Junos OS Release 14.1X53-D30 for QFX Series switches. **auto option** was also added at this time.

auto option added in Junos OS Release 19.1R1 for MX series.

Statement introduced in cRPD Release 19.4R1.

Description

Specify a virtual routing and forwarding (VRF) target community. If you configure the **community** option only, default VRF import and export policies are generated that accept and tag routes with the specified target community. The purpose of the **vrf-target** statement is to simplify the configuration by allowing you to configure most statements at the **[edit routing-instances]** hierarchy level. In effect, this statement configures a single policy for import and a single policy for export to replace the per-VRF policies for every community.

You can still create more complex policies by explicitly configuring VRF import and export policies using the **import** and **export** options.

Options

community—Community name.

auto—Automatically derives the route target (RT). The auto-derived route targets have higher precedence over manually configured RT in vrf-target, vrf-export policies, and vrf-import policies.

NOTE: Auto-derived route targets are supported only in virtual switch and EVPN routing instances.

import community-name—Allowed communities accepted from neighbors.

export community-name—Allowed communities sent to neighbors.

Required Privilege Level

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Policies for the VRF Table on PE Routers in VPNs

Example: Configuring FEC 129 BGP Autodiscovery for VPWS

Operational Commands

IN THIS CHAPTER

- `clear ancp access-loop` | 510
- `clear ancp neighbor` | 512
- `clear dhcp relay binding` | 514
- `clear dhcp relay statistics` | 518
- `clear dhcp server binding` | 521
- `clear dhcp server statistics` | 525
- `clear dhcpv6 server binding` | 527
- `clear dhcpv6 server statistics` | 530
- `clear network-access aaa subscriber` | 532
- `request ancp oam port-down` | 534
- `request ancp oam port-up` | 536
- `request auto-configuration reconnect-pending` | 538
- `show ancp neighbor` | 539
- `show ancp subscriber` | 549
- `show auto-configuration out-of-band` | 559
- `show dhcp relay binding` | 565
- `show dhcp relay statistics` | 572
- `show dhcp server binding` | 577
- `show dhcp server statistics` | 587
- `show dhcpv6 server binding` | 592
- `show dhcpv6 server statistics` | 600
- `show extensible-subscriber-services counters` | 604
- `show extensible-subscriber-services debug-information` | 606
- `show extensible-subscriber-services dictionary` | 607
- `show extensible-subscriber-services dictionary attributes` | 612
- `show extensible-subscriber-services dictionary services` | 615
- `show extensible-subscriber-services sessions` | 619
- `show extensible-subscriber-services service` | 621

- [show interfaces \(Fast Ethernet\) | 623](#)
- [show interfaces | 646](#)
- [show interfaces \(Loopback\) | 753](#)
- [show interfaces \(PPPoE\) | 763](#)
- [show interfaces demux0 \(Demux Interfaces\) | 776](#)
- [show interfaces filters | 788](#)
- [show interfaces l2-routing-instance | 791](#)
- [show interfaces routing | 793](#)
- [show interfaces routing-instance | 801](#)
- [show network-access aaa statistics | 804](#)
- [show network-access aaa statistics authentication | 817](#)
- [show network-access aaa subscribers | 821](#)
- [show network-access address-assignment pool | 827](#)
- [show ppp interface | 829](#)
- [show subscribers | 849](#)
- [show subscribers summary | 896](#)
- [show vpls connections | 905](#)
- [show vpls flood event-queue | 922](#)
- [show vpls flood instance | 924](#)
- [show vpls flood route | 927](#)
- [show vpls mac-table | 930](#)
- [show vpls statistics | 937](#)

clear ancp access-loop

Syntax

```
clear ancp access-loop
(neighbor ip-address | subscriber-interface physical-interface-name)
circuit-id aci
remote-id ari
outer-vlan-id vlan-id
```

Release Information

Command introduced in Junos OS Release 16.1R4.

Description

Clear the connection for the subscriber on the specified access loop for an ANCP-triggered, autosensed dynamic VLAN. The autoconfiguration daemon (autoconfd) deletes any existing cached information about the subscriber. This command simulates a CPE connection reset as seen when the access node sends a Port Down message followed by a Port Up message.

Options

aci—ANCP Access-Loop-Circuit-ID TLV that identifies the subscriber-side access loop logical port and partially identifies an access loop to clear.

ari—ANCP Access-Loop-Remote-ID TLV that uniquely identifies the subscriber on the access loop and partially identifies an access loop to clear.

ip-address—ANCP neighbor's IP address that specifies an access loop to clear.

physical-interface-name—Subscriber interface that specifies an access loop to clear.

vlan-id—ANCP Access-Aggregation-Circuit-ID-Binary TLV that identifies the logical circuit identifier on the NAS side and partially identifies an access loop to clear.

Required Privilege Level

clear

RELATED DOCUMENTATION

[Clearing ANCP Access Loops | 147](#)

[Triggering ANCP OAM to Simulate ANCP Port Down and Port Up Messages | 142](#)

[Layer 2 Wholesale with ANCP-Triggered VLANs Overview | 116](#)

List of Sample Output

[clear ancp access-loop on page 511](#)

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the **show ancp neighbor** command before and after clearing the access loop to verify the clear operation.

Sample Output

```
clear ancp access-loop
```

```
user@host> clear ancp-access-loop neighbor 192.168.25.31 circuit-id line-aci-1 remote-id line-ari-1  
outer-vlan-id 126
```

clear ancp neighbor

Syntax

```
clear ancp neighbor  
<ip-address ip-address>  
<system-name mac-address>
```

Release Information

Command introduced in Junos OS Release 9.4.

Description

Clear the ANCP agent connection with all ANCP neighbors or with the specified ANCP neighbor. This command deletes information for subscribers associated with the neighbor, causing the adjusted traffic rates to revert to the configured rate for the subscriber interfaces. The neighbor remains configured (its administrative state is *enabled*) and can reestablish adjacencies.

This command initiates logout of ANCP-triggered dynamic VLAN sessions on the physical interface associated with the specified neighbor; conventionally autosensed dynamic VLAN sessions and their associated logical interfaces are not affected.

Options

none—Clear all ANCP neighbors.

ip-address *ip-address*—(Optional) Clear the ANCP neighbor specified by the IP address.

system-name *mac-address*—(Optional) Clear the ANCP neighbor specified by the MAC address.

Required Privilege Level

clear

RELATED DOCUMENTATION

| [show ancp neighbor](#) | 539

List of Sample Output

[clear ancp neighbor on page 513](#)

[show ancp neighbor on page 513](#)

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the **show ancp neighbor** command before and after clearing the ANCP neighbors to verify the clear operation.

Sample Output

clear ancp neighbor

```
user@host> clear ancp neighbor
```

show ancp neighbor

The following sample output displays the connections with ANCP neighbors before and after the **clear ancp neighbor** command was issued.

```
user@host> show ancp neighbor
```

IP Address	MAC Address	State	Subscriber Count	Capabilities
203.0.113.102	00:00:5e:00:53:10	Established	5	Topo
203.0.113.122	00:00:5e:00:53:12	Established	5	Topo
203.0.113.132	00:00:5e:00:53:13	Established	5	Topo
203.0.113.142	00:00:5e:00:53:14	Established	5	Topo

```
user@host> clear ancp neighbor ip-address 203.0.113.102
```

```
user@host> show ancp neighbor
```

IP Address	MAC Address	State	Subscriber Count	Capabilities
203.0.113.122	00:00:5e:00:53:12	Established	5	Topo
203.0.113.132	00:00:5e:00:53:13	Established	5	Topo
203.0.113.142	00:00:5e:00:53:14	Established	5	Topo

clear dhcp relay binding

Syntax

```
clear dhcp relay binding
<address>
<all>
<dual-stack>
<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.

Option **dual-stack** added in Junos OS Release 15.1.

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.

dual-stack—(Optional) Clear the binding state for DHCPv4 clients and the associated DHCPv6 bindings in the single-session DHCP dual stack. DHCPv6 clients created in a DHCPv6-only stack are not affected.

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

Viewing and Clearing DHCP Bindings

[show dhcp relay binding](#) | [565](#)

List of Sample Output

- [clear dhcp relay binding on page 515](#)
- [clear dhcp relay binding all on page 516](#)
- [clear dhcp relay binding dual-stack all on page 516](#)
- [clear dhcp relay binding interface on page 516](#)
- [clear dhcp relay binding <interfaces-vlan> on page 516](#)
- [clear dhcp relay binding <interfaces-wildcard> on page 516](#)

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
198.51.100.32	00:00:5e:00:53:01	active	2007-02-08 16:41:17 EST
192.168.14.8	00:00:5e:00:53:02	active	2007-02-10 10:01:06 EST

user@host> **clear dhcp relay binding 198.51.100.32**

user@host> **show dhcp relay binding**

IP address	Hardware address	Type	Lease expires at
192.168.14.8	00:00:5e:00:53:02	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 dual-stack all

The following command clears all DHCP relay agent bindings for all DHCPv4 clients and the associated DHCPv6 bindings in the single-session DHCP dual stack. DHCPv6 clients created in a DHCPv6-only stack are not affected.

```
user@host> clear dhcp relay binding dual-stack 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
<bulk-leasequery-connections>
<leasequery>
<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.

bulk-leasequery-connections option introduced in Junos OS Release 16.1.

Description

Clear all Dynamic Host Configuration Protocol (DHCP) relay statistics.

Options

bulk-leasequery-connections—(Optional) Clear DHCP relay bulk leasequery statistics.

leasequery—(Optional) Clear DHCP relay individual leasequery statistics.

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](#) | 572

List of Sample Output

[clear dhcp relay statistics on page 519](#)

Output Fields

See [show dhcp relay statistics](#) for an explanation of output fields.

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                1
  Lease Time Violated  1

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

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>
<dual-stack>
```

Release Information

Command introduced in Junos OS Release 9.0.

Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

Command updated with **dual-stack** statement in Junos OS Release 17.3.

Description

Clear the binding state of a Dynamic Host Configuration Protocol (DHCP) client from the client table on the extended DHCP local server.

NOTE: If you delete the DHCP server configuration, DHCP server bindings might still remain. To ensure that DHCP bindings are removed, issue the **clear dhcp server binding** command before you delete the DHCP server configuration.

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.

dual-stack—(Optional) Remove either both arms or single arm of dual-stack.

NOTE:

- The **dual-stack** command is added in the syntax removes both arms of the dual-stack with a single command entry.
- When the **dual-stack** command is not added in the syntax, the **clear dhcpv6 server binding** command clears only the family specific arm of the dual-stack.

Required Privilege Level

view

RELATED DOCUMENTATION

Viewing and Clearing DHCP Bindings

[show dhcp server binding](#) | 577

List of Sample Output

[clear dhcp server binding <ip-address> on page 523](#)

[clear dhcp server binding all on page 523](#)

[clear dhcp server binding interface on page 524](#)

[clear dhcp server binding <interfaces-vlan> on page 524](#)

[clear dhcp server binding <interfaces-wildcard> on page 524](#)

[clear dhcp server binding dual-stack all on page 524](#)

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
198.51.100.1    00:00:5e:00:53:01 active    2007-01-17 11:38:47 PST
198.51.100.3    00:00:5e:00:53:02 active    2007-01-17 11:38:41 PST
```

user@host> **clear dhcp server binding 198.51.100.1**

user@host> **show dhcp server binding**

```
1 clients, (0 bound, 0 selecting, 0 renewing, 0 rebinding)

IP address      Hardware address  Type      Lease expires at
198.51.100.3    00:00:5e:00:53:02 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 binding dual-stack all**

The following command clears all the dual-stack local server bindings.

```
user@host> clear dhcp server binding dual-stack all
```


clear dhcp server statistics

Syntax

```
clear dhcp server statistics  
<bulk-leasequery-connections>  
<logical-system logical-system-name>  
<routing-instance routing-instance-name>
```

Release Information

Command introduced in Junos OS Release 9.0.

bulk-leasequery-connections option introduced in Junos OS Release 16.1.

Description

Clear all extended Dynamic Host Configuration Protocol (DHCP) local server statistics.

Options

bulk-leasequery-connections—(Optional) Clear DHCPv4 local server bulk leasequery statistics.

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 525](#)

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                1
    Lease Time Violation 1

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

```

clear dhcpv6 server binding

Syntax

```
clear dhcpv6 server binding
<address>
<all>
<interface interface-name>
<interfaces-vlan>
<interfaces-wildcard>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
<dual-stack>
```

Release Information

Command introduced in Junos OS Release 9.6.

Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

Command updated with **dual-stack** statement in Junos OS Release 17.3.

Description

Clear the binding state of a Dynamic Host Configuration Protocol for IPv6 (DHCPv6) client from the client table on the extended DHCPv6 local server.

Options

address—(Optional) Clear the binding state for the DHCPv6 client, using one of the following entries:

- *CID*—The specified Client ID (CID).
- *ipv6-prefix*—The specified IPv6 prefix.
- *session-id*—The specified session ID.

all—(Optional) Clear the binding state for all DHCPv6 clients.

interface *interface-name*—(Optional) Clear the binding state for DHCPv6 clients on the specified interface.

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 DHCPv6 clients on the specified logical system.

routing-instance *routing-instance-name*—(Optional) Clear the binding state for DHCPv6 clients on the specified routing instance.

dual-stack—(Optional) Remove either both arms or single arm of dual-stack.

NOTE:

- The **dual-stack** command is added in the syntax removes both arms of the dual-stack with a single command entry.
- When the **dual-stack** command is not added in the syntax, the **clear dhcpv6 server binding** command clears only the family specific arm of the dual-stack.

Required Privilege Level

clear

RELATED DOCUMENTATION

Viewing and Clearing DHCP Bindings

[show dhcpv6 server binding](#) | 592

List of Sample Output

[clear dhcpv6 server binding all on page 528](#)

[clear dhcpv6 server binding <ipv6-prefix> on page 529](#)

[clear dhcpv6 server binding interface on page 529](#)

[clear dhcpv6 server binding <interfaces-vlan> on page 529](#)

[clear dhcpv6 server binding <interfaces-wildcard> on page 529](#)

[clear dhcpv6 server binding dual-stack all on page 529](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear dhcpv6 server binding all

The following command clears all DHCPv6 local server bindings:

```
user@host> clear dhcpv6 server binding all
```

clear dhcpv6 server binding <ipv6-prefix>

The following command clears DHCPv6 local server bindings for a specific IPv6 prefix:

```
user@host> clear dhcpv6 server binding 14/0x00010001/0x02b3be8f/0x00109400/0x0005
```

clear dhcpv6 server binding interface

The following command clears DHCPv6 local server bindings on a specific interface:

```
user@host> clear dhcpv6 server binding interface fe-0/0/2
```

clear dhcpv6 server binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCPv6 local server bindings on top of the underlying interface **ae0**, which clears DHCPv6 bindings on all demux VLANs on top of **ae0**:

```
user@host> clear dhcpv6 server binding interface ae0
```

clear dhcpv6 server binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCPv6 local server bindings over a specific interface:

```
user@host> clear dhcpv6 server binding ge-1/0/0.*
```

clear dhcpv6 server binding dual-stack all

The following command clears all the dual-stack local server bindings.

```
user@host> clear dhcpv6 server binding dual-stack all
```

clear dhcpv6 server statistics

Syntax

```
clear dhcpv6 server statistics  
<bulk-leasequery-connections>  
<interface interface-name>  
<logical-system logical-system-name>  
<routing-instance routing-instance-name>
```

Release Information

Command introduced in Junos OS Release 9.6.

bulk-leasequery-connections option introduced in Junos OS Release 16.1.

Description

Clear all extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server statistics.

Options

bulk-leasequery-connections—(Optional) Clear DHCPv6 local server bulk leasequery statistics.

logical-system *logical-system-name*—(Optional) Clear the statistics for DHCPv6 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 DHCPv6 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

clear

RELATED DOCUMENTATION

[show dhcpv6 server statistics](#) | 600

List of Sample Output

[clear dhcpv6 server statistics on page 531](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

```
clear dhcpv6 server statistics
```

```
user@host> clear dhcpv6 server statistics
```

clear network-access aaa subscriber

Syntax

```
clear network-access aaa subscriber
<session-id identifier <reconnect>>
<statistics username username>
<username username <reconnect>>
```

Release Information

Command introduced in Junos OS Release 9.1.

reconnect and **session-id** options added in Junos OS Release 16.1R4.

Description

Clear AAA subscriber statistics and log out subscribers. You can log out subscribers based on the username or on the subscriber session identifier. Use the session identifier when more than one session has the same username string.

Options

reconnect—(Optional) Reconnect as a Layer 2 wholesale session when the subscriber session has been fully logged out. This option is equivalent to issuing a RADIUS-initiated disconnect with reconnect semantics; that is, when the message includes Acct-Terminate-Cause (RADIUS attribute 49) with a value of callback (16). You can apply this option to either a Layer 2 wholesale session or a conventionally auto-sensed dynamic VLAN supporting a PPPoE session.

In the latter case, this option triggers a PPPoE session logout and removal of the dynamic VLAN logical interface. This is followed by authorization of the access-line to attempt creation of a dynamic VLAN IFL supporting Layer 2 wholesale session in its place.

session-id *identifier*—(Optional) Log out the subscriber based on the subscriber session identifier.

statistics username *username*—(Optional) Clear AAA subscriber statistics and log out the subscriber.

username *username*—(Optional) Log out the AAA subscriber.

Required Privilege Level

maintenance

RELATED DOCUMENTATION

Verifying and Managing Subscriber AAA Information

List of Sample Output

[clear network-access aaa subscriber statistics username on page 533](#)

[clear network-access aaa subscriber username on page 533](#)

[clear network-access aaa subscriber session-id on page 533](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear network-access aaa subscriber statistics username

```
user@host> clear network-access aaa subscriber statistics username user22@example.com
```

clear network-access aaa subscriber username

```
user@host> clear network-access aaa subscriber username user22@example.com
```

clear network-access aaa subscriber session-id

```
user@host> clear network-access aaa subscriber session-id 18367425
```

request ancp oam port-down

Syntax

```
request ancp oam port-down
(neighbor ip-address | subscriber-interface physical-interface-name)
circuit-id aci remote-id ari outer-vlan-id vlan-id
```

Release Information

Command introduced in Junos OS Release 16.1R4.

Description

Simulate an ANCP Port Down message on the specified access loop for troubleshooting or to mitigate an abnormal condition. Triggers removal of the corresponding out-of-band triggered, autosensed dynamic VLAN session for which no ANCP-sourced information exists. You must specify an ACI, an ARI, and an outer VLAN tag. This command is overridden by a genuine ANCP Port-Up message, meaning that you cannot use this command to initiate a Port Down condition when the access node has already reported a Port Up condition.

Options

aci—ANCP Access-Loop-Circuit-ID TLV that corresponds to a subscriber interface on the access node; used to identify the access node from which the message is simulated.

ari—ANCP Access-Loop-Remote-ID TLV that identifies the subscriber associated with an interface on the access node; used to identify the access node from which the message is simulated.

ip-address—IP address that specifies the access node from which the message is simulated.

physical-interface-name—Name of the access-facing subscriber interface that specifies the access node on whose local loop the loopback test is run.

vlan-id—ANCP Access-Aggregation-Circuit-ID-Binary TLV, the outer VLAN tag inserted by the access node on upstream traffic; used to identify the access node from which the message is simulated.

Required Privilege Level

view

RELATED DOCUMENTATION

[Triggering ANCP OAM to Simulate ANCP Port Down and Port Up Messages | 142](#)

[Layer 2 Wholesale with ANCP-Triggered VLANs Overview | 116](#)

List of Sample Output

[request ancp oam port-down neighbor circuit-id remote-id outer-vlan-id on page 535](#)

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the **show ancp neighbor detail**, **show subscribers client-type vlan-oob detail**, and the **show subscribers summary** commands before and after initiating the Port Down message to verify the operation.

Sample Output

```
request ancp oam port-down neighbor circuit-id remote-id outer-vlan-id
```

```
user@host> request ancp oam port-down neighbor 192.168.25.31 circuit-id line-aci-1 remote-id  
line-ari-1 outer-vlan-id 126
```

request ancp oam port-up

Syntax

```
request ancp oam port-up
(neighbor ip-address | subscriber-interface physical-interface-name)
circuit-id aci remote-id ari outer-vlan-id vlan-id
```

Release Information

Command introduced in Junos OS Release 16.1R4.

Description

Simulate an ANCP Port Up message on the specified access loop for troubleshooting or to mitigate an abnormal condition. You must specify an ACI, an ARI, and an outer VLAN tag. This command is overridden by a genuine ANCP Port Down message, meaning that you cannot use this command to initiate a Port Up condition when the access node has already reported a Port Down condition.

Options

aci—ANCP Access-Loop-Circuit-ID TLV that corresponds to a subscriber interface on the access node; used to identify the access node from which the message is simulated.

ip-address—IP address that specifies the access node from which the message is simulated.

ari—ANCP Access-Loop-Remote-ID TLV that identifies the subscriber associated with an interface on the access node; used to identify the access node from which the message is simulated.

physical-interface-name—Name of the access-facing subscriber interface that specifies the access node on whose local loop the loopback test is run.

vlan-id—ANCP Access-Aggregation-Circuit-ID-Binary TLV, the outer VLAN tag inserted by the access node on upstream traffic; used to identify the access node from which the message is simulated.

Required Privilege Level

view

RELATED DOCUMENTATION

[Triggering ANCP OAM to Simulate ANCP Port Down and Port Up Messages | 142](#)

[Layer 2 Wholesale with ANCP-Triggered VLANs Overview | 116](#)

List of Sample Output

[request ancp oam port-up neighbor circuit-id remote-id outer-vlan-id on page 537](#)

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the **show ancp neighbor detail**, **show subscribers client-type vlan-oob detail**, and the **show subscribers summary** commands before and after initiating the Port Up message to verify the operation.

Sample Output

```
request ancp oam port-up neighbor circuit-id remote-id outer-vlan-id
```

```
user@host> request ancp oam port-up neighbor 192.168.25.31 circuit-id line-aci-1 remote-id line-ari-1  
outer-vlan-id 126
```

request auto-configuration reconnect-pending

Syntax

```
request auto-configuration reconnect-pending
```

Release Information

Command introduced in Junos OS Release 16.1R4.

Description

Initiate reestablishment of Layer 2 wholesale sessions that correspond to access lines that are in the pending state. Ordinarily, the most likely situations with pending sessions are handled automatically. This statement is intended to be used only when an uncommon condition might prevent automatic reestablishment. This command has no effect when no pending sessions are present.

Required Privilege Level

view

RELATED DOCUMENTATION

[Reestablishing Pending Access Line Sessions for Layer 2 Wholesale | 145](#)

[Layer 2 Wholesale with ANCP-Triggered VLANs Overview | 116](#)

List of Sample Output

[request auto-configuration reconnect-pending on page 538](#)

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the **show ancp neighbor detail**, **show subscribers client-type vlan-oob detail**, and the **show subscribers summary** commands before and after initiating the Port Up message to verify the operation.

Sample Output

```
request auto-configuration reconnect-pending
```

```
user@host> request auto-configuration reconnect-pending
```

show ancp neighbor

Syntax

```
show ancp neighbor
<brief | detail>
<ip-address ip-address
<system-name mac-address>
```

Release Information

Command introduced in Junos OS Release 9.4.

Description

Display information about all ANCP neighbors or the specified ANCP neighbor, regardless of operational state.

Options

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

ip-address ip-address —(Optional) Display information about the neighbor (access node) specified by the IP address.

system-name mac-address—(Optional) Display information about the neighbor (access node) specified by the MAC address.

Required Privilege Level

view

RELATED DOCUMENTATION

show ancp cos
show ancp subscriber 549

List of Sample Output

- [show ancp neighbor on page 544](#)
- [show ancp neighbor detail on page 544](#)
- [show ancp neighbor ip-address on page 546](#)
- [show ancp neighbor system-name on page 547](#)

Output Fields

[Table 10 on page 540](#) lists the output fields for the **show ancp neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 10: show ancp neighbor Output Fields

Field Name	Field Description	Level of Output
Version	Version of the ANCP implementation: <ul style="list-style-type: none"> • 0x31—General Switch Management Protocol (GSMP) version 3, sub-version 1; ANCP version before <i>RFC 6320, Protocol for Access Node Control Mechanism in Broadband Networks</i>. • 0x32—ANCP version 1, defined in <i>RFC 6320, Protocol for Access Node Control Mechanism in Broadband Networks</i>. 	brief detail none
IP Address	IP address of the ANCP neighbor.	brief detail none
PartId	Number that associates the ANCP message with a specific partition.	brief none
State	Operational state of the ANCP adjacency: <ul style="list-style-type: none"> • Configured—The neighbor has been configured, but has never been in the Established state. An asterisk (*) is prefixed to the neighbor entry for this state. • Establishing—Adjacency negotiations are in progress for the neighbor. An asterisk (*) is prefixed to the neighbor entry for this state. This state is rarely seen because the adjacency is established so quickly. • Established—Adjacency negotiations have succeeded for the neighbor and an ANCP session has been established. • Not Estblshed—Not Established; adjacency negotiations are ready to begin. Indicates that this neighbor previously had been in the Established state; that is, it has lost a previously established adjacency. An asterisk (*) is prefixed to the neighbor entry for this state. 	All levels
Time	How long the adjacency has been up in one of the following formats: <ul style="list-style-type: none"> • <i>nwndnh</i>—number of weeks, days, and hours • <i>nd hh:mm:ss</i>—number of days, hours, minutes, and seconds 	brief detail none
Subscriber Count	Number of subscribers associated with the ANCP neighbor (access local loop).	brief none
Capabilities	Negotiated ANCP capability: <ul style="list-style-type: none"> • Topo—Topology discovery. • OAM—Performance of local Operations Administration Maintenance (OAM) procedures on an access loop controlled by the router. 	All levels

Table 10: show ancp neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
System Name	MAC address of the ANCP neighbor.	detail
TCP Port	TCP port on which ANCP messages are exchanged.	detail
System Instance	Number identifying the ANCP link instance from the edge device's perspective.	detail
Peer Instance	Number identifying the ANCP instance from the access node's perspective. This number is unique and changes when the node or link comes back up after going down.	detail
Timer	Adjacency timer value advertised by the ANCP peer in 100 ms increments; the interval between ANCP ACK messages. This value remains constant for the duration of an ANCP session.	detail
Partition Type	Number that identifies whether partitions are used and how the ID is negotiated: <ul style="list-style-type: none"> • 0—No partition. • 1—Fixed partition requested. • 2—Fixed partition assigned. 	detail
Partition Flag	Number that specifies the type of partition requested: 1 (new adjacency) or 2 (recovered adjacency).	detail
Partition Identifier	Number that identifies a logical partition of an access node with which the ANCP agent has formed an adjacency. A value of zero indicates that the agent supports each neighbor on an IP address over a single TCP session with a partition ID of zero. This is the default support case. A nonzero value indicates that the agent supports each neighbor on an IP address over a single TCP session with a nonzero partition ID.	detail
Partition Adjacencies	Number of adjacencies that share the partition.	detail
Dead Timer	Remaining period that the edge device waits for adjacency packets from a neighbor before declaring the neighbor to be down. The maximum dead time value is three times the configured adjacency timer value. This field displays the current value based on the time that the last adjacency packet was received.	detail

Table 10: show ancp neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
Received Syn Count	Number of synchronization messages received from neighbors to maintain adjacencies.	detail
Received Synack Count	Number of synchronization acknowledgment messages received from neighbors in response to the node's synchronization messages.	detail
Received Rstack Count	Number of messages received from neighbors indicating that the link to the neighbor needs to be reset.	detail
Received Ack Count	Number of acknowledgment messages periodically received from neighbors after an adjacency has been established.	detail
Received Port Up Count	Number of status messages received from neighbors indicating that a port has transitioned to the up state.	detail
Received Port Down Count	Number of status messages received from neighbors indicating that a port has transitioned to the down state.	detail
Received Generic Resp Count	Number of generic response messages received from neighbors.	detail
Received Adjacency Update Count	Number of adjacency update messages received from neighbors.	detail
Received OAM Count	Number of OAM responses received from neighbors in reply to request commands.	detail
Received Other Count	Number of all other ANCP message packets received from neighbors that do not fit into one of the other categories.	detail
Sent Syn Count	Number of synchronization messages sent to neighbors to maintain adjacencies.	detail
Sent Synack Count	Number of synchronization acknowledgment messages sent to neighbors in response to the their synchronization messages.	detail
Sent Rstack Count	Number of messages sent to neighbors indicating that the link to the neighbor needs to be reset.	detail

Table 10: show ancp neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
Sent Ack Count	Number of acknowledgment messages periodically sent to neighbors after an adjacency has been established.	detail
Sent Generic Resp Count	Number of generic response messages sent to neighbors.	detail
Sent OAM Count	Number of OAM request commands sent to neighbors.	detail
Max Discovery Limit Exceed Count	Number of times that the maximum number of discovery table entries accepted from the neighbor has been exceeded.	detail
Result Codes	<p>Number of generic response messages sent to neighbors that include each of the following result codes:</p> <ul style="list-style-type: none"> • Invalid Request Message Count—A properly formed request message violated the protocol because of timing (such as a race condition) or direction of transmission. • Specified Port(s) Down Count—One or more of the specified ports are down because of a state mismatch between the router and an ANCP control application. • Out of Resources Count—ANCP is out of resources, probably not related to the access lines. This result code is sent only by an access node. • Request Msg Not Implemented Count— • Malformed Msg Count—Message is malformed because it was corrupted in transit or there was an implementation error at either end of the connection. • TLV Missing Count—One or more mandatory TLVs was missing from a request. • Invalid TLV Contents Count—The contents of one or more TLVs in the request do not match its required specification. • Non-Existent Port(s) Count—One or more of the ports specified in a request do not exist, possibly because of a configuration mismatch between the access node and the router or AAA. 	detail

Sample Output

show ancp neighbor

```
user@host> show ancp neighbor
```

Version	IP Address	PartID	State	Time	Subscriber
Capabilities					Count
0x31	203.0.113.13	0	Established	11:24	2
Topo					
0x31	203.0.113.15	0	Not Estblshd	2:45	2
Topo					
* 0x0	198.51.100.102	0	Establishing	0	0
* 0x0	192.0.2.0	0	Configured	0	0
* 0x0	192.0.2.1	0	Configured	0	0

show ancp neighbor detail

```
user@host> show ancp neighbor detail
```

```
Neighbor Information
  Version           : 0x31
  IP Address        : 192.0.2.85
  System Name       : 00:00:5e:00:53:01
    Up Time          : 26
    TCP Port         : 32666
    State            : Established
    Subscriber Count : 4
    Capabilities     : Topo
    System Instance  : 2
    Peer Instance    : 20
    Adjacency Timer (in 100ms) : 100
    Peer Adjacency Timer (in 100ms) : 100
    Partition Type   : 0
    Partition Flag    : 1
    Partition Identifier : 0
    Partition Adjacencies : 0
    Dead Timer       : 23
    Received Syn Count : 1
    Received Synack Count : 1
    Received Rstack Count : 0
    Received Ack Count : 4
    Received Port Up Count : 10
```

Received Port Down Count	: 0	
Received Generic Resp Count	: 0	
Received Adjacency Update Count	: 0	
Received OAM Count	: 0	
Received Other Count	: 0	
Sent Syn Count	: 1	
Sent Synack Count	: 2	
Sent Rstack Count	: 0	
Sent Ack Count	: 3	
Sent Generic Resp Count	: 0	
Sent OAM Count	: 0	
Max Discovery Limit Exceed Count	: 0	
Result Codes:	Received	Sent
Invalid Request Message Count	: 0	0
Specified Port(s) Down Count	: 0	0
Out of Resources Count	: 0	0
Request Msg Not Implemented Count	: 0	0
Malformed Msg Count	: 0	0
TLV Missing Count	: 0	0
Invalid TLV Contents Count	: 0	0
Non-Existent Port(s) Count	: 0	0

Version	: 0x32
IP Address	: 192.168.9.1
System Name	: 00:00:5e:00:53:02
Up Time	: 36
TCP Port	: 61408
State	: Not Established
Subscriber Count	: 1
Capabilities	: Topology Discovery
System Instance	: 12
Peer Instance	: 1
Adjacency Timer (in 100ms)	: 50
Peer Adjacency Timer (in 100ms)	: 100
Partition Type	: 0
Partition Flag	: 1
Partition Identifier	: 0
Partition Adjacencies	: 0
Dead Timer	: 23
Received Syn Count	: 24
Received Synack Count	: 20
Received Rstack Count	: 2
Received Ack Count	: 9

```

Received Port Up Count      : 5
Received Port Down Count    : 0
Received Generic Resp Count : 0
Received Adjacency Update Count : 0
Received OAM Responses Count : 2
Received Other Count        : 0
Sent Syn Count              : 20
Sent Synack Count           : 24
Sent Rstack Count           : 1
Sent Generic Resp Count     : 0
Sent Ack Count              : 9
Sent OAM Requests Count     : 4
Max Discovery Limit Exceed Count : 0
Result Codes:
Invalid Request Message Count : 0
Specified Port(s) Down Count  : 0
Out of Resources Count         : 0
Request Msg Not Implemented Count: 0
Malformed Msg Count           : 0
TLV Missing Count             : 0
Invalid TLV Contents Count     : 0
Non-Existent Port(s) Count     : 0
Received                      Sent

```

show ancp neighbor ip-address

user@host> show ancp neighbor ip-address 192.0.2.85

```

Neighbor Information
Version           : 0x32
IP Address        : 192.0.2.85
System Name       : 00:00:5e:00:53:ba
Up Time           : 26
TCP Port          : 32666
State             : Established
Subscriber Count  : 4
Capabilities      : Topo
System Instance   : 2
Peer Instance     : 20
Adjacency Timer (in 100ms) : 100
Peer Adjacency Timer (in 100ms) : 100
Partition Type    : 0
Partition Flag    : 1
Partition Identifier : 0

```

```

Partition Adjacencies           : 0
Dead Timer                      : 23
Received Syn Count              : 1
Received Synack Count           : 1
Received Rstack Count           : 0
Received Ack Count              : 4
Received Port Up Count          : 10
Received Port Down Count        : 0
Received Generic Resp Count     : 0
Received Adjacency Update Count : 0
Received OAM Count              : 0
Received Other Count            : 0
Sent Syn Count                  : 1
Sent Synack Count               : 2
Sent Rstack Count               : 0
Sent Ack Count                  : 3
Sent Generic Resp Count         : 0
Sent OAM Count                  : 0
Max Discovery Limit Exceed Count : 0
Result Codes:
Invalid Request Message Count   : 0
Specified Port(s) Down Count    : 0
Out of Resources Count          : 0
Request Msg Not Implemented Count : 0
Malformed Msg Count             : 0
TLV Missing Count               : 0
Invalid TLV Contents Count      : 0
Non-Existent Port(s) Count      : 0
Received                        : 0
Sent                           : 0

```

show ancp neighbor system-name

user@host> show ancp neighbor 00:00:5e:00:53:ba detail

```

Neighbor Information
Version           : 0x31
IP Address        : 203.0.113.101
System Name       : 00:00:5e:00:53:ba
Up Time           : 19
TCP Port          : 1028
State             : Established
Subscriber Count  : 2
Capabilities      : Topology Discovery, OAM
System Instance   : 1

```

Peer Instance	:	10	
Adjacency Timer (in 100ms)	:	100	
Peer Adjacency Timer (in 100ms)	:	250	
Partition Type	:	0	
Partition Flag	:	1	
Partition Identifier	:	0	
Partition Adjacencies	:	0	
Dead Timer	:	55	
Received Syn Count	:	1	
Received Synack Count	:	1	
Received Rstack Count	:	0	
Received Ack Count	:	1	
Received Port Up Count	:	34	
Received Port Down Count	:	0	
Received Generic Resp Count	:	0	
Received Adjacency Update Count	:	0	
Received OAM Responses Count	:	2	
Received Other Count	:	0	
Sent Syn Count	:	1	
Sent Synack Count	:	1	
Sent Rstack Count	:	0	
Sent Ack Count	:	3	
Sent Generic Resp Count	:	0	
Sent OAM Requests Count	:	4	
Max Discovery Limit Exceed Count	:	3	
Result Codes:	Received		Sent
Invalid Request Message Count	:	0	0
Specified Port(s) Down Count	:	0	0
Out of Resources Count	:	0	0
Request Msg Not Implemented Count	:	0	0
Malformed Msg Count	:	0	0
TLV Missing Count	:	0	0
Invalid TLV Contents Count	:	0	0
Non-Existent Port(s) Count	:	0	0

show ancp subscriber

Syntax

```
show ancp subscriber
<brief | detail>
<access-aggregation-circuit-id circuit-identifier>
<identifier identifier>
<ip-address ip-address>
<system-name mac-address>
```

Release Information

Command introduced in Junos OS Release 9.4.

neighbor option replaced with **ip-address** in Junos OS Release 16.1.

system-name option introduced in Junos OS Release 16.1.

access-aggregation-circuit-id option introduced in Junos OS Release 18.4R1.

Description

Display information about active subscribers regardless of the subscriber's operational state, for all subscribers (local access loops), the subscriber associated with the access line specified by an ACI, or the subscriber associated with the specified ANCP neighbor (access node).

After an ancpd restart, this command displays orphaned entries (marked with an **o**) for subscriber sessions that were established before the restart but which have not yet been reestablished. As sessions are reestablished, the number of orphaned entries displayed by the command decreases. The number reaches zero when all sessions are reestablished or when the orphaned-interface timer expires.

Options

none—Display information about all subscribers.

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

access-aggregation-circuit-id *circuit-identifier*—(Optional) Display information about ANCP subscribers whose Access-Aggregation-Circuit-ID-ASCII attribute (TLV 0x0003) matches the specified value.

A *circuit-identifier* that begins with the **#** character indicates a backhaul line identifier. You can specify a wildcard (*) anywhere in the string.

identifier *identifier*—(Optional) Display information about the subscriber associated with the access line (ACI) specified by the access identifier.

ip-address *ip-address* —(Optional) Display information about the subscribers connected to the access node specified by the IP address.

system-name mac-address—(Optional) Display information about the subscribers connected to the access node specified by the MAC address.

Required Privilege Level

view

RELATED DOCUMENTATION

<i>clear ancp subscriber</i>
<i>show ancp cos</i>
show ancp neighbor 539
<i>show ancp statistics</i>

List of Sample Output

- [show ancp subscriber on page 554](#)
- [show ancp subscriber \(After ancpd Restart\) on page 555](#)
- [show ancp subscriber brief on page 555](#)
- [show ancp subscriber detail on page 555](#)
- [show ancp subscriber access-aggregation-circuit-id detail on page 557](#)
- [show ancp subscriber identifier identifier-string detail on page 557](#)

Output Fields

[Table 11 on page 551](#) lists the output fields for the **show ancp subscriber** command. Output fields are listed in the approximate order in which they appear.

Table 11: show ancp subscriber Output Fields

Field Name	Field Description	Level of Output
Loop Identifier	<p>Access loop identifier as sent by the access node and configured to map the subscriber to an interface.</p> <p>An asterisk (*) indicates that the information might be stale due to receiving a Port Down message with a DSL Line State of Idle.</p> <p>Two asterisks (**) indicate that the neighbor associated with the subscriber has lost its adjacency. In this case, the DSL Line State might be Established.</p> <p>An o indicates that the entry is for an orphaned interface and represents a previously established subscriber session that has not been reestablished after an ancpc restart.</p> <p>The number of orphaned entries decreases as the ANCP neighbors reestablish adjacencies and the protocol subscriber sessions are reestablished. The command output indicates this by removing the o marker.</p> <p>Eventually the number of orphaned entries reaches zero, because either all the adjacencies and subscriber sessions have been reestablished or any remaining orphaned entries are removed when the orphaned-interface timer expires.</p>	brief none
DSL Line State	State of the DSL line: Idle , Showtime , or Silent .	brief detail
Access Type	Type of access line employed by the access node: ADSL1 , ADSL2 , ADSL2+ , VDSL1 , VDSL2 , SDSL , G.fast , VDSL2 Annex Q , SDSL bonded , VDSL2 bonded , G.fast bonded , VDSL2 Annex Q bonded or OTHER .	brief detail none
Interface	Name of the interface set or logical interface.	brief detail none
Rate Kbps	Actual downstream data rate for this local loop.	brief none
Neighbor	IP address of ANCP neighbor (access node).	brief none

Table 11: show ancp subscriber Output Fields (*continued*)

Field Name	Field Description	Level of Output
Access Loop Circuit Identifier	<p>Access loop circuit identifier as sent by the access node and configured to map the subscriber to an interface.</p> <p>An asterisk (*) indicates that the information might be stale due to receiving a Port Down message with a DSL Line State of Idle.</p> <p>Two asterisks (**) indicate that the neighbor associated with the subscriber has lost its adjacency. In this case, the DSL Line State might be Established.</p>	detail
Neighbor IP Address	IP address of the ANCP neighbor (access node).	detail
Aggregate Circuit Identifier	ASCII identifier for the subscriber access loop; value of the Access-Aggregation-Circuit-ID-ASCII attribute (TLV 0x0003).	detail
Aggregate Circuit Identifier Binary	Binary identifier for the VLAN circuit ID.	detail
Tech Type	Type of technology employed by the subscriber. Currently Junos OS supports DSL technology type only.	detail
DSL Line Data Link	Data link protocol employed on the access loop: AAL5 or Ethernet .	detail
DSL Line Encapsulation	<p>Encapsulation type on the access loop, for Ethernet only:</p> <ul style="list-style-type: none"> • 0—NA, type not conveyed • 1—Untagged Ethernet • 2—Single-tagged Ethernet 	detail

Table 11: show ancp subscriber Output Fields (*continued*)

Field Name	Field Description	Level of Output
DSL Line Encapsulation Payload	Payload carried across the access loop: <ul style="list-style-type: none"> • 0—NA, type not conveyed • 1—PPPoA LLC • 2—PPPoA null • 3—IPoA LLC • 4—IPoA null • 5—Ethernet over AAL5 LLC with FCS • 6—Ethernet over AAL5 LLC without FCS • 7—Ethernet over AAL5 null with FCS • 8—Ethernet over AAL5 null without FCS 	detail
Interface Type	Type of interface employed for subscriber traffic: ifl for a single VLAN or interface-set for a configured group of VLANs.	detail
Actual Net Data Upstream	Actual upstream data rate for this local loop, in Kbps.	detail
Actual Net Data Downstream	Actual downstream data rate for this local loop, in Kbps.	detail
Minimum Net Data Upstream	Minimum upstream data rate desired by the operator for this local loop, in Kbps.	detail
Minimum Net Data Downstream	Minimum downstream data rate desired by the operator for this local loop, in Kbps.	detail
Maximum Net Data Upstream	Maximum upstream data rate desired by the operator for this local loop, in Kbps.	detail
Maximum Net Data Downstream	Maximum downstream data rate desired by the operator for this local loop, in Kbps.	detail
Attainable Net Data Upstream	Maximum attainable upstream data rate for this local loop, in Kbps.	detail

Table 11: show ancp subscriber Output Fields (*continued*)

Field Name	Field Description	Level of Output
Attainable Net Data Downstream	Maximum attainable downstream data rate for this local loop, in Kbps.	detail
Minimum Low Power Data Downstream	Minimum downstream data rate desired by the operator for this local loop in low power state, in Kbps.	detail
Minimum Low Power Data Upstream	Minimum upstream data rate desired by the operator for this local loop in low power state, in Kbps.	detail
Maximum Interleave Delay Downstream	Maximum interleaving delay for downstream data, in milliseconds.	detail
Maximum Interleave Delay Upstream	Maximum interleaving delay for upstream data, in milliseconds.	detail
Actual Interleave Delay Downstream	Actual interleaving delay for downstream data, in milliseconds.	detail
Actual Interleave Delay Upstream	Actual interleaving delay for upstream data, in milliseconds.	detail

Sample Output

show ancp subscriber

user@host> **show ancp subscriber**

Loop Identifier Rate Neighbor	DSL Line State	Tech Type	Access Type	Interface Kbps
**circuit 101 203.0.113.13	Idle	DSL	ADSL1	---- 32
**circuit 102	Idle	DSL	ADSL1	---- 32

```

203.0.113.13
  circuit 301          Showtime DSL      ADSL1    ----          32
203.0.113.15
  circuit 302          Showtime DSL      ADSL1    ----          32
203.0.113.15

```

show ancp subscriber (After ancpd Restart)

```
user@host> show ancp subscriber
```

```

  Loop Identifier      DSL Line Tech Type      Access Type Interface
  Rate      Neighbor
  State
o circuit 201          Showtime DSL      ADSL1    ----          222222
o circuit 202          Showtime DSL      ADSL1    ----          222222

```

show ancp subscriber brief

```
user@host> show ancp subscriber brief
```

```

  Loop Identifier      Type      Interface      Rate      Neighbor
                                Kbps
port-1-10             VDSL2     set-ge-10410   64        203.0.113.102
port-1-11             VDSL2     set-ge-10411   64        203.0.113.111
port-2-10             VDSL2     ge-1/0/4.12    64        203.0.113.112
port-2-11             VDSL2     ge-1/0/4.13    64        203.0.113.113

```

show ancp subscriber detail

```
user@host> show ancp subscriber detail
```

```

Subscriber Information
* Access Loop Circuit Identifier : circuit 101
  Neighbor IP Address           : 203.0.113.13
  Aggregate Circuit Identifier Binary : 0/0
  Tech Type                     : DSL

```

```

Access Type                : ADSL1
DSL Line State              : Idle
DSL Line Data Link          : Data link 2
DSL Line Encapsulation      : N/A
DSL Line Encapsulation Payload : N/A
Interface Type              : N/A
Interface                   : ----
Actual Net Data Upstream    : 32
Actual Net Data Downstream  : 32
Minimum Net Data Upstream   : 0
Minimum Net Data Downstream : 0
Maximum Net Data Upstream   : 0
Maximum Net Data Downstream : 0
Attainable Net Data Upstream : 1024
Attainable Net Data Downstream : 8192
Minimum Low Power Data Downstream : 32
Minimum Low Power Data Upstream : 32
Maximum Interleave Delay Downstream : 20
Maximum Interleave Delay Upstream : 20
Actual Interleave Delay Downstream : 20
Actual Interleave Delay Upstream : 20
* Access Loop Circuit Identifier: circuit 102
  Neighbor IP Address        : 213.0.113.13
  Aggregate Circuit Identifier Binary : 0/0
  Tech Type                  : DSL
  Access Type                : ADSL1
  DSL Line State              : Idle
  DSL Line Data Link          : Data link 2
  DSL Line Encapsulation      : N/A
  DSL Line Encapsulation Payload : N/A
  Interface Type              : N/A
  Interface                   : ----
  Actual Net Data Upstream    : 32
  Actual Net Data Downstream  : 32
  Minimum Net Data Upstream   : 0
  Minimum Net Data Downstream : 0
  Maximum Net Data Upstream   : 0
  Maximum Net Data Downstream : 0
  Attainable Net Data Upstream : 1024
  Attainable Net Data Downstream : 8192
  Minimum Low Power Data Downstream : 32
  Minimum Low Power Data Upstream : 32
  Maximum Interleave Delay Downstream : 20
  Maximum Interleave Delay Upstream : 20

```



```

Actual Interleave Delay Downstream : 20
Actual Interleave Delay Upstream   : 20
...
```

show ancp subscriber access-aggregation-circuit-id detail

user@host> show ancp subscriber access-aggregation-circuit-id "#TEST-DPU-C-100" detail

```

Subscriber Information
* Access Loop Circuit Identifier : circuit 201
  Neighbor IP Address             : 192.0.2.1
  Access Loop Remote Identifier   : remote 123
  Aggregate Circuit Identifier    : #TEST-DPU-C-100
  Aggregate Circuit Identifier Binary : 50
  Tech Type:                      : DSL
  Interface Type                  : interface
  Interface                       : ge-1/0/0.3221225475
  Actual Net Data Upstream        : 1024
  Actual Net Data Downstream      : 2048
  Maximum Net Data Upstream       : 0
  Maximum Net Data Downstream     : 0
* Access Loop Circuit Identifier : circuit 202
  Neighbor IP Address             : 192.0.2.1
  Access Loop Remote Identifier   : remote 185
  Aggregate Circuit Identifier    : #TEST-DPU-C-100
  Aggregate Circuit Identifier Binary : 50
  Tech Type:                      : DSL
  Interface Type                  : interface
  Interface                       : ge-1/0/0.3221225476
  Actual Net Data Upstream        : 1024
  Actual Net Data Downstream      : 2048
  Maximum Net Data Upstream       : 0
  Maximum Net Data Downstream     : 0
```

show ancp subscriber identifier identifier-string detail

user@host> show ancp subscriber identifier port-1-11 detail

```

Access Loop Identifier : port-1-11
  Neighbor IP Address   : 203.0.113.112
  Aggregate Circuit Identifier Binary : 0/0
```

DSL Type	: DSL 0
Interface Type	: interface-set
Interface	: set-ge-10411
DSL Line State	: Show Time
Actual Net Data Upstream	: 64
Actual Net Data Downstream	: 64
DSL Line Data Link	: AAL5
DSL Line Encapsulation	: N/A
DSL Line Encapsulation Payload	: N/A
Minimum Net Data Upstream	: 64
Minimum Net Data Downstream	: 64
Maximum Net Data Upstream	: 64
Maximum Net Data Downstream	: 64
Attainable Net Data Upstream	: 64
Attainable Net Data Downstream	: 64
Minimum Low Power Data Downstream	: 64
Minimum Low Power Data Upstream	: 64
Maximum Interleave Delay Downstream	: 50
Maximum Interleave Delay Upstream	: 50
Actual Interleave Delay Downstream	: 50
Actual Interleave Delay Upstream	: 50

show auto-configuration out-of-band

Syntax

```
show auto-configuration out-of-band  
< brief | count | detail >  
< all | pending >  
< circuit-id >  
< interface >
```

Release Information

Command introduced in Junos OS Release 16.1R4.

all, **circuit-id**, **count**, and **interface** options introduced in Junos OS Release 17.3R1.

Description

Display information about ANCP-triggered out-of-band subscriber sessions. You can display all such subscribers or subscribers associated with a particular ACL or interface. You can also display information for sessions that are in the pending state. Sessions transition to the pending state when the Layer 2 wholesale session is authorized and assigned to an existing, nondefault routing instance, but subsequently profile instantiation to create the dynamic VLAN logical interface fails.

Options

brief | count | detail—(Optional) Display the specified level of output. The default output level is the same as **brief**. The **count** level displays the number of out-of-band subscribers and sessions.

all—(Optional) Display information for all out-of-band subscriber sessions, including pending sessions.

circuit-id—(Optional) Limit the display to subscriber sessions on the specified access loop circuit identifier.

interface—(Optional) Limit the display to out-of-band subscriber sessions on the specified interface.

pending—(Optional) Display access lines in the pending state per routing instance.

Required Privilege Level

view

RELATED DOCUMENTATION

[Layer 2 Wholesale with ANCP-Triggered VLANs Overview](#) | 116

List of Sample Output

[show auto-configuration out-of-band on page 561](#)

[show auto-configuration out-of-band \(All Sessions\) on page 561](#)

[show auto-configuration out-of-band \(All Sessions, Detail\) on page 562](#)

[show auto-configuration out-of-band \(Brief\) on page 562](#)
[show auto-configuration out-of-band \(Circuit ID\) on page 563](#)
[show auto-configuration out-of-band \(Circuit ID, Detail\) on page 563](#)
[show auto-configuration out-of-band \(Count\) on page 563](#)
[show auto-configuration out-of-band \(Detail\) on page 563](#)
[show auto-configuration out-of-band \(Interface\) on page 564](#)
[show auto-configuration out-of-band \(Interface, Detail\) on page 564](#)
[show auto-configuration out-of-band \(Pending Sessions\) on page 564](#)

Output Fields

Table 12 on page 560 lists the output fields for the **show auto-configuration out-of-band** command.

Table 12: show auto-configuration out-of-band Output Fields

Field Name	Field Description	Level of Output
Agent Circuit ID	Agent circuit ID (ACI, access loop circuit identifier) that maps the subscriber to an interface.	detail
Agent Remote ID	Agent remote ID (ARI, access loop remote identifier) that identifies the subscriber associated with an interface and therefore the access node.	detail
Circuit Id	Agent circuit ID (ACI, access loop circuit identifier) that maps the subscriber to an interface.	brief none
Core IFL name	Name of the core-facing logical interface.	detail
Inner Vlan Map Id	Inner VLAN tag that replaces the outer VLAN tag when the subscriber traffic is tunneled to the network service provider router.	detail
Login Time	Timestamp when the subscriber logged in.	detail
Out-of-band subscribers	Number of active out-of-band subscribers.	brief none
Pending count	Number of subscriber sessions in the routing instance that are in the pending state.	brief none
Remote Id	Agent remote ID (ARI, access loop remote identifier) for the subscriber.	brief none
Routing-Instance	Name of the routing instance for which the pending count is displayed.	brief none

Table 12: show auto-configuration out-of-band Output Fields (*continued*)

Field Name	Field Description	Level of Output
Session ID	Subscriber session identifier.	All levels
Subscriber IFL	Name of the subscriber logical interface.	brief none
Subscriber Interface	Name of the subscriber logical interface.	detail
total entries	Total number of active plus pending subscriber sessions.	brief none
VLAN	VLAN ID of the subscriber' s incoming interface.	brief none
VLAN ID	VLAN ID of the subscriber' s incoming interface.	detail
Vlan Map Id	Trunk VLAN ID assigned to the core-facing physical interface within its input VLAN map.	detail

Sample Output

show auto-configuration out-of-band

```
user@host> show auto-configuration out-of-band
```

Circuit Id ID VLAN	Remote Id	Subscriber IFL	Session
203.0.13.3/0.0.0.0 eth 1/0 200	arn.1211.000000000002	ge-2/1/1.1689443721	1301

show auto-configuration out-of-band (All Sessions)

```
user@host> show auto-configuration out-of-band all
```

Circuit Id ID VLAN	Remote Id	Subscriber IFL	Session
203.0.13.2/0.0.0.0 eth 1/0 100	arn.4312.000000000002	ge-2/0/4.2547532212	1287
203.0.13.3/0.0.0.0 eth 1/0 200	arn.1211.000000000002	ge-2/1/1.1689443721	1301

```
203.0.13.4/0.0.0.0 eth 1/0  arn.5628.000000000002  ge-2/1/2.7697428544  1628
300
```

show auto-configuration out-of-band (All Sessions, Detail)

user@host> show auto-configuration out-of-band all detail

```
Agent Circuit ID: 203.0.13.2/0.0.0.0 eth 1/0
Agent Remote ID: arn.4312.000000000002
Subscriber interface: ge-2/0/4.2547532212
Session ID: 1287
VLAN ID: 100
Core IFL name: ae1.0
Vlan Map Id: 2
Inner Vlan Map Id: 1010
Login Time: 2018-06-25 09:05:14 EST
Agent Circuit ID: 203.0.13.3/0.0.0.0 eth 1/0
Agent Remote ID: arn.1211.000000000002
Subscriber interface: ge-2/1/1.1689443721
Session ID: 1301
VLAN ID: 200
Core IFL name: ae1.1
Vlan Map Id: 3
Inner Vlan Map Id: 1020
Login Time: 2018-06-25 09:11:21 EST
Agent Circuit ID: 203.0.13.4/0.0.0.0 eth 1/0
Agent Remote ID: arn.5628.000000000002
Subscriber interface: ge-2/1/2.7697428544
Session ID: 1628
VLAN ID: 300
Core IFL name: ae1.2
Vlan Map Id: 4
Inner Vlan Map Id: 1030
Login Time: 2018-06-25 09:11:58 EST
```

show auto-configuration out-of-band (Brief)

user@host> show auto-configuration out-of-band brief

Circuit Id ID VLAN	Remote Id	Subscriber IFL	Session
203.0.13.3/0.0.0.0 eth 1/0 200	arn.1211.000000000002	ge-2/1/1.1689443721	1301

show auto-configuration out-of-band (Circuit ID)

```
user@host> show auto-configuration out-of-band circuit-id "203.0.13.2/0.0.0.0 eth 1/0"
```

Circuit Id ID VLAN	Remote Id	Subscriber IFL	Session
203.0.13.2/0.0.0.0 eth 1/0 100	arn.4312.000000000002	ge-2/0/4.2547532212	1287

show auto-configuration out-of-band (Circuit ID, Detail)

```
user@host> show auto-configuration out-of-band circuit-id "203.0.13.2/0.0.0.0 eth 1/0" detail
```

```
Agent Circuit ID: 203.0.13.2/0.0.0.0 eth 1/0
Agent Remote ID: arn.4312.000000000002
Subscriber interface: ge-2/0/4.2547532212
Session ID: 1287
VLAN ID: 100
Core IFL name: ae1.0
Vlan Map Id: 2
Inner Vlan Map Id: 1010
Login Time: 2018-06-25 09:05:14 EST
```

show auto-configuration out-of-band (Count)

```
user@host> show auto-configuration out-of-band count
```

```
Out-of-band subscribers: 1, total entries: 1
```

show auto-configuration out-of-band (Detail)

```
user@host> show auto-configuration out-of-band detail
```

```
Agent Circuit ID: ACI.0001.aa1000000111
Agent Remote ID: remote-id
Subscriber interface: ge-1/0/0.3221225472
Session ID: 1
VLAN ID: 51
Core IFL name: ge-1/0/4.0
Vlan Map Id: 20
Inner Vlan Map Id: 1
Login Time: 2018-07-10 10:27:55 EDT
```

show auto-configuration out-of-band (Interface)

```
user@host> show auto-configuration out-of-band interface ge-2/0/4.2547532212
```

Circuit Id ID VLAN	Remote Id	Subscriber IFL	Session
203.0.13.2/0.0.0.0 eth 1/0 100	arn.4312.000000000002	ge-2/0/4.2547532212	1287

show auto-configuration out-of-band (Interface, Detail)

```
user@host> show auto-configuration out-of-band interface ge-2/0/4.2547532212 detail
```

```
Agent Circuit ID: 203.0.13.2/0.0.0.0 eth 1/0
Agent Remote ID: arn.4312.000000000002
Subscriber interface: ge-2/0/4.2547532212
Session ID: 1287
VLAN ID: 100
Core IFL name: ae1.0
Vlan Map Id: 2
Inner Vlan Map Id: 1010
Login Time: 2018-06-25 09:05:14 EST
```

show auto-configuration out-of-band (Pending Sessions)

```
user@host> show auto-configuration out-of-band pending
```

Circuit Id ID VLAN	Remote Id	Subscriber IFL	Session
203.0.13.2/0.0.0.0 eth 1/0 100	arn.4312.000000000002	ge-2/0/4.2547532212	

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

Viewing and Clearing DHCP Bindings clear dhcp relay binding 514
--

List of Sample Output

- [show dhcp relay binding on page 568](#)
- [show dhcp relay binding detail on page 568](#)
- [show dhcp relay binding interface on page 569](#)
- [show dhcp relay binding interface vlan-id on page 569](#)
- [show dhcp relay binding interface svlan-id on page 570](#)
- [show dhcp relay binding ip-address on page 570](#)
- [show dhcp relay binding mac-address on page 570](#)
- [show dhcp relay binding session-id on page 570](#)
- [show dhcp relay binding <interfaces-vlan> on page 570](#)
- [show dhcp relay binding <interfaces-wildcard> on page 571](#)
- [show dhcp relay binding summary on page 571](#)

Output Fields

Table 13 on page 566 lists the output fields for the **show dhcp relay binding** command. Output fields are listed in the approximate order in which they appear.

Table 13: show dhcp relay binding Output Fields

Field Name	Field Description	Level of Output
<i>number</i> clients,(<i>number</i> init, <i>number</i> bound, <i>number</i> selecting, <i>number</i> requesting, <i>number</i> renewing, <i>number</i> rebinding, <i>number</i> releasing)	Summary counts of the total number of DHCP clients and the number of DHCP clients in each state.	summary

Table 13: show dhcp relay binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
IP address	IP address of the DHCP client.	brief detail
Session Id	Session ID of the subscriber session.	brief detail
Generated Remote ID	Remote ID generated by the Option 82 Agent Remote ID (suboption 1)	detail
Hardware address	Hardware address of the DHCP client.	brief detail
Expires	Number of seconds in which the lease expires.	brief detail
State	<p>State of the DHCP relay address binding table on the DHCP client:</p> <ul style="list-style-type: none"> • 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. 	brief detail
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
Lease time violated	Lease time violation has occurred.	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

Table 13: show dhcp relay binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
Bootp Relay Address	IP address of BOOTP relay.	detail
Type	Type of DHCP packet processing performed on the router: <ul style="list-style-type: none"> • 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
Dual Stack Group	Name of dual stack that is configured with the DHCP binding.	detail
Dual Stack Peer Prefix	Prefix of dual stack DHCPv6 peer.	detail
Dual Stack Peer Address	Address of the dual stack DHCPv6 peer.	detail

Sample Output

show dhcp relay binding

```
user@host> show dhcp relay binding
```

IP address	Session Id	Hardware address	Expires	State	Interface
198.51.100.11	41	00:00:5e:00:53:01	86371	BOUND	ge-1/0/0.0
198.51.100.12	42	00:00:5e:00:53:02	86371	BOUND	ge-1/0/0.0
198.51.100.13	43	00:00:5e:00:53:03	86371	BOUND	ge-1/0/0.0
198.51.100.14	44	00:00:5e:00:53:04	86371	BOUND	ge-1/0/0.0
198.51.100.15	45	00:00:5e:00:53:05	86371	BOUND	ge-1/0/0.0

show dhcp relay binding detail

```
user@host> show dhcp relay binding detail
```

```

Client IP Address: 198.51.100.11
  Hardware Address: 00:00:5e:00:53:01
  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
  Lease time violated: yes
  Last Packet Received: 2009-07-20 11:00:06 PDT
  Incoming Client Interface: ge-1/0/0.0
  Server Ip Address: 198.51.100.22
  Server Interface: none
  Bootp Relay Address: 198.51.100.32
  Session Id: 41
  Dual Stack Group: dual-stack-retail6
  Dual Stack Peer Prefix: 2001:db8:0:4::/64
  Dual Stack Peer Address: 2001:db8:1:0:8003::1/128

```

```

Client IP Address: 198.51.100.12
  Hardware Address: 00:00:5e:00:53:02
  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: 198.51.100.22
  Server Interface: none
  Bootp Relay Address: 198.51.100.32
  Session Id: 42
  Generated Remote ID: host:ge-1/0/0:100

```

show dhcp relay binding interface

```
user@host> show dhcp relay binding interface fe-0/0/2
```

IP address	Hardware address	Type	Lease expires at
198.51.100.1	00:00:5e:00:53:01	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
198.51.100.15	6	00:00:5e:00:53:94	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
198.51.100.16	7	00:00:5e:00:53:92	86124	BOUND	ge-1/1/0:10-100

show dhcp relay binding ip-address

user@host> show dhcp relay binding 198.51.100.13

IP address	Session Id	Hardware address	Expires	State	Interface
198.51.100.13	43	00:00:5e:00:53:03	86293	BOUND	ge-1/0/0.0

show dhcp relay binding mac-address

user@host> show dhcp relay binding 00:00:5e:00:53:05

IP address	Session Id	Hardware address	Expires	State	Interface
198.51.100.15	45	00:00:5e:00:53:05	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
198.51.100.11	41	00:00:5e:00:53:53	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	00:00:5e:00:53:02	86346	BOUND	ge-1/0/0.1073741827
192.168.0.16	41	00:00:5e:00:53:01	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	00:00:5e:00:53:04	86361	BOUND	ge-1/3/0.110
192.168.0.8	23	00:00:5e:00:53:03	86361	BOUND	ge-1/3/0.110
192.168.0.7	22	00:00:5e:00:53:02	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
<bulk-leasequery-connections>
<leasequery>
<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.

bulk-leasequery-connections option introduced in Junos OS Release 16.1.

Description

Display Dynamic Host Configuration Protocol (DHCP) relay statistics.

Options

bulk-leasequery-connections—(Optional) Display information about DHCP relay bulk leasequery statistics.

leasequery—(Optional) Display information about DHCP relay individual leasequery statistics.

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](#) | 518

List of Sample Output

[show dhcp relay statistics on page 575](#)

[show dhcp relay statistics bulk-leasequery-connections on page 576](#)

Output Fields

[Table 14 on page 573](#) lists the output fields for the **show dhcp relay statistics** command. Output fields are listed in the approximate order in which they appear.

Table 14: show dhcp relay statistics Output Fields

Field Name	Field Description
Packets dropped	<p>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.</p> <ul style="list-style-type: none"> • 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. • Lease Time Violation—Number of packets discarded because of a lease time violation • 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.

Table 14: show dhcp relay statistics Output Fields (*continued*)

Field Name	Field Description
Messages received	<p>Number of DHCP messages received.</p> <ul style="list-style-type: none"> • 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 • DHCPLEASEACTIVE—Number of active DHCP leases • DHCPLEASEUNASSIGNED—Number of DHCP leases that are managed by the server but have not yet been assigned • DHCPLEASEUNKNOWN—Number of unknown DHCP leases • DHCPLEASEQUERYDONE—The leasequery is complete
Messages sent	<p>Number of DHCP messages sent.</p> <ul style="list-style-type: none"> • 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 • DHCPLEASEQUERY—Number of DHCP leasequery messages transmitted • DHCPLEASEBULKLEASEQUERY—Number of DHCP bulk leasequery messages transmitted
External Server Response	State of the external DHCP server responsiveness.
Packets forwarded	<p>Number of packets forwarded.</p> <ul style="list-style-type: none"> • BOOTREQUEST—Number of BOOTREQUEST protocol data units (PDUs) forwarded • BOOTREPLY—Number of BOOTREPLY protocol data units (PDUs) forwarded
External Server Response	State of the external DHCP server responsiveness.
Total Requested Servers	Total number of servers with which the DHCP relay agent has requested a bulk leasequery connection.
Total Attempted Servers	Total number of servers with which the DHCP relay agent has attempted to create a bulk leasequery connection.

Table 14: show dhcp relay statistics Output Fields (*continued*)

Field Name	Field Description
Total Connected	Total number of servers that have formed a bulk leasequery connection with the DHCP relay agent.
Total Terminated by Server	Total number of servers that have terminated a bulk leasequery connection with the DHCP relay agent.
Total Max Attempted	Total number of servers where the DHCP relay agent reached the maximum retry limit when it attempted to create a bulk leasequery connection.
Total Closed due to Errors	Total number of bulk leasequery connections that closed due to an internal error on the DHCP relay agent.
In-Flight Connected	Number of current bulk leasequery connections on the DHCP relay agent.
Bulk Leasequery Reply Packet Retries	Number of bulk leasequery reply packets that the DHCP relay agent has retried.

Sample Output

show dhcp relay statistics

user@host> **show dhcp relay statistics**

```

Packets dropped:
  Total                34
  Bad hardware address  1
  Bad opcode           1
  Bad options          3
  Invalid server address 5
  Lease Time Violation 1
  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
    DHCPLEASEACTIVE       0
    DHCPLEASEUNASSIGNED   0
    DHCPLEASEUNKNOWN      0
    DHCPLEASEQUERYDONE    0

```

```

Messages sent:
    BOOTREPLY             0
    DHCPOFFER              2
    DHCPACK                1
    DHCPNAK                0
    DHCPFORCERENEW         0
    DHCPLEASEQUERY         0
    DHCPBULKLEASEQUERY     0

```

```

Packets forwarded:
    Total                  4
    BOOTREQUEST            2
    BOOTREPLY              2

```

```

External Server Response:
    State                  Responding

```

show dhcp relay statistics bulk-leasequery-connections

```
user@host> show dhcp relay statistics bulk-leasequery-connections
```

```

Total Requested Servers:    0
Total Attempted Servers:   0
Total Connected:            0
Total Terminated by Server: 0
Total Max Attempted:        0
Total Closed due to Errors: 0
In-Flight Connected:        0
Bulk Leasequery Reply Packet Retries: 0

```

show dhcp server binding

Syntax

```
show dhcp server binding
<address>
<interfaces-vlan><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.

NOTE: If you delete the DHCP server configuration, DHCP server bindings might still remain. To ensure that DHCP bindings are removed, issue the **clear dhcp server binding** command before you delete the DHCP server configuration.

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

Viewing and Clearing DHCP Bindings
Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration
clear dhcp server binding 521

List of Sample Output

- [show dhcp server binding on page 582](#)
- [show dhcp server binding detail on page 582](#)
- [show dhcp server binding detail \(ACI Interface Set Configured\) on page 584](#)
- [show dhcp server binding interface <vlan-id> on page 584](#)
- [show dhcp server binding interface <svlan-id> on page 584](#)
- [show dhcp server binding <ip-address> on page 585](#)
- [show dhcp server binding <session-id> on page 585](#)
- [show dhcp server binding summary on page 585](#)
- [show dhcp server binding <interfaces-vlan> on page 585](#)
- [show dhcp server binding <interfaces-wildcard> on page 585](#)

Output Fields

[Table 15 on page 579](#) lists the output fields for the **show dhcp server binding** command. Output fields are listed in the approximate order in which they appear.

Table 15: show dhcp server binding Output Fields

Field Name	Field Description	Level of Output
<i>number</i> clients, (<i>number</i> init, <i>number</i> bound, <i>number</i> selecting, <i>number</i> requesting, <i>number</i> renewing, <i>number</i> 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: <ul style="list-style-type: none"> • 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
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

Table 15: show dhcp server binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
Lease Start	Date and time at which the client's IP address lease started.	detail
Lease time violated	Lease time violation has occurred.	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

Table 15: show dhcp server binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
Liveness Detection State	<p>State of the liveness detection status for a subscriber's Bidirectional Forwarding Detection (BFD) protocol session:</p> <p>NOTE: This output field displays status only when liveness detection has been explicitly configured for a subscriber and the liveness detection protocol is actively functioning for that subscriber.</p> <ul style="list-style-type: none"> • DOWN—Liveness detection has been enabled for a subscriber but the broadband network gateway (BNG) detects that the liveness detection session for the BFD protocol is in the DOWN state. A liveness detection session that was previously in an UP state has transitioned to a DOWN state, beginning with a liveness detection failure, and ending with the deletion of the client binding. The DOWN state is reported only during this transition period of time. • UNKNOWN—Liveness detection has been enabled for a subscriber but the actual liveness detection state has not yet been determined. The UNKNOWN state is reported after a DHCP subscriber initially logs in while the underlying liveness detection protocol handshake, such as BFD, is still processing and the BFD session has not yet reached the UP state. • UP—Liveness detection has been enabled for a subscriber, and the BNG and the subscriber or client have <i>both</i> determined that the liveness detection session for the BFD protocol is in the UP state. • WENT_DOWN—State is functionally equivalent to the DOWN state. A liveness detection session that was previously in an UP state has transitioned to a DOWN state implying a liveness detection failure. The WENT_DOWN state applies to the internal distribution of the liveness detection mechanism between the Junos DHCP Daemon for Subscriber Services (JDHCPd), the BFD plug-in within the Broadband Edge Subscriber Management Daemon (BBE-SMGD), and the Packet Forwarding Engine. 	detail

Table 15: show dhcp server binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
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
Client Profile Name	DHCP client profile name.	detail
Dual Stack Group	DHCP server profile name.	detail
Dual Stack Peer Prefix	IPv6 prefix of peer.	detail
Dual Stack Peer Address	IPv6 address of peer.	detail

Sample Output

show dhcp server binding

```
user@host> show dhcp server binding
```

IP address	Session Id	Hardware address	Expires	State	Interface
198.51.100.15	6	00:00:5e:00:53:01	86180	BOUND	ge-1/0/0.0
198.51.100.16	7	00:00:5e:00:53:02	86180	BOUND	ge-1/0/0.0
198.51.100.17	8	00:00:5e:00:53:03	86180	BOUND	ge-1/0/0.0
198.51.100.18	9	00:00:5e:00:53:04	86180	BOUND	ge-1/0/0.0
198.51.100.19	10	00:00:5e:00:53:05	86180	BOUND	ge-1/0/0.0

show dhcp server binding detail

```
user@host> show dhcp server binding detail
```

```

Client IP Address: 198.51.100.15
  Hardware Address:      00:00:5e:00:53:01
  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:     198.51.100.9
  Server Interface:      none
  Session Id:            6
  Client Pool Name:      6
  Liveness Detection State: UP

Client IP Address:      198.51.100.16
  Hardware Address:      00:00:5e:00:53:02
  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
  Lease time violated:    yes
  Incoming Client Interface: ge-1/0/0.0
  Server Ip Address:     198.51.100.9
  Server Interface:      none
  Session Id:            7
  Client Pool Name:      7
  Liveness Detection State: UP

```

When DHCP binding is configured with dual-stack, we get the following output:

user@host> **show dhcp server binding detail**

```

Client IP Address: 198.51.100.10
  Hardware Address:      00:00:64:03:01:02
  State:                 BOUND(LOCAL_SERVER_STATE_BOUND)
  Protocol-Used:         DHCP
  Lease Expires:         2016-11-07 08:30:39 PST
  Lease Expires in:      43706 seconds
  Lease Start:           2016-11-04 11:00:37 PDT
  Last Packet Received:  2016-11-06 09:00:39 PST
  Incoming Client Interface: ae0.3221225472
  Client Interface Svlan Id: 2000
  Client Interface Vlan Id: 1
  Server Ip Address:     198.51.100.2

```

```

Session Id:                2
Client Pool Name:          my-v4-pool
Client Profile Name:       dhcp-retail
Dual Stack Group:          my-dual-stack
Dual Stack Peer Prefix:    2001:db8:ffff:0:4::/64
Dual Stack Peer Address:   2001:db8:0:8003::1/128

```

show dhcp server binding detail (ACI Interface Set Configured)

user@host> show dhcp server binding detail

```

Client IP Address: 198.51.100.14
  Hardware Address: 00:00:5e:00:53:02
  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: 198.51.100.202
  Session Id:        11
  Client Pool Name:   poolA
  Client Profile Name: DEMUXprofile
  Liveness Detection State: UP
  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
198.51.100.15	6	00:00:5e:00:53:01	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
198.51.100.16	7	00:00:5e:00:53:02	86124	BOUND	ge-1/1/0:10-100

show dhcp server binding <ip-address>

user@host> show dhcp server binding 198100.19

IP address	Session Id	Hardware address	Expires	State	Interface
198.51.100.19	10	00:00:5e:00:53:05	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
198.51.100.15	6	00:00:5e:00:53:01	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	00:00:5e:00:53:02	86346	BOUND	ge-1/0/0.1073741827
192.168.0.16	41	00:00:5e:00:53:01	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	00:00:5e:00:53:04	86361	BOUND	ge-1/3/0.110

192.168.0.8	23	00:00:5e:00:53:03	86361	BOUND	ge-1/3/0.110
192.168.0.7	22	00:00:5e:00:53:02	86361	BOUND	ge-1/3/0.110

show dhcp server statistics

Syntax

```
show dhcp server statistics
<bulk-leasequery-connections>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Release Information

Command introduced in Junos OS Release 9.0.

bulk-leasequery-connections option introduced in Junos OS Release 16.1.

Description

Display extended Dynamic Host Configuration Protocol (DHCP) local server statistics.

Options

bulk-leasequery-connections—(Optional) Display information about DHCP local server bulk leasequery statistics.

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](#) | 525

List of Sample Output

[show dhcp server statistics on page 590](#)

[show dhcp server statistics on page 590](#)

Output Fields

[Table 16 on page 588](#) lists the output fields for the **show dhcp server statistics** command. Output fields are listed in the approximate order in which they appear.

Table 16: show dhcp server statistics Output Fields

Field Name	Field Description
Packets dropped	<p>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.</p> <ul style="list-style-type: none"> • 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 • Lease Time Violation—Number of packets discarded because of a lease time violation • 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
Offer Delay	<p>Number of DHCPv4 offer messages delayed.</p> <ul style="list-style-type: none"> • DELAYED—Number of DHCPv4 offer packets that have been sent after being delayed. • INPROGRESS—Number of DHCPv4 offer packets that are in the delay queue. • TOTAL—Total number of delayed DHCPv4 offer messages; sum of DELAYED and INPROGRESS.

Table 16: show dhcp server statistics Output Fields (*continued*)

Field Name	Field Description
Messages received	<p>Number of DHCP messages received.</p> <ul style="list-style-type: none"> • 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 • DHCPLEASEQUERY—Number of DHCP leasequery messages received. • DHCPBULKLEASEQUERY—Number of DHCP bulk leasequery messages received. • DHCPRENEW—Number of DHCP renew messages received; subset of DHCPREQUEST counter. • DHCPREBIND—Number of DHCP rebind messages received; subset of DHCPREQUEST counter.
Messages sent	<p>Number of DHCP messages sent.</p> <ul style="list-style-type: none"> • 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 • DHCPLEASEUNASSIGNED—Number of DHCP leases that are managed by the server but have not yet been assigned • DHCPLEASEUNKNOWN—Number of unknown DHCP leases • DHCPLEASEACTIVE—Number of active DHCP leases • DHCPLEASEQUERYDONE—The leasequery is complete
Total Accepted Connections	Total number of bulk leasequery connections accepted by the server.
Total Not-Accepted Connections	Total number of bulk leasequery connections not accepted by the server.
Connections Closed due to Errors	Number of bulk leasequery connections that the server closed due to an internal error.
Connections Closed due to max-empty-replies	Number of bulk leasequery connections that the server closed because the maximum number of empty replies was reached.

Table 16: show dhcp server statistics Output Fields (*continued*)

Field Name	Field Description
In-flight Connections	Number of bulk leasequery connections on the server.

Sample Output

show dhcp server statistics

user@host> show dhcp server statistics

```

Packets dropped:
    Total                1
    Lease Time Violation 1

Offer Delay:
    DELAYED              3
    INPROGRESS            9
    TOTAL                 12

Messages received:
    BOOTREQUEST          25
    DHCPDECLINE           0
    DHCPDISCOVER          10
    DHCPINFORM            0
    DHCPRELEASE           4
    DHCPREQUEST           10
    DHCPRENEW              4
    DHCPREBIND            2

Messages sent:
    BOOTREPLY             20
    DHCPOFFER             10
    DHCPACK                10
    DHCPNAK                0
    DHCPFORCERENEW         0

```

show dhcp server statistics

user@host> show dhcp server statistics verbose

Packets dropped:	
Total	0
Messages received:	
BOOTREQUEST	238
DHCPDECLINE	0
DHCPDISCOVER	1
DHCPINFORM	0
DHCPRELEASE	0
DHCPREQUEST	237
DHCPRENEW	236
DHCPREBIND	0
Messages sent:	
BOOTREPLY	20
DHCPOFFER	10
DHCPACK	10
DHCPNAK	0
DHCPFORCERENEW	0

show dhcpv6 server binding

Syntax

```
show dhcpv6 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 9.6.

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 for IPv6 (DHCPv6) local server.

Options

address—(Optional) One of the following identifiers for the DHCPv6 client whose binding state you want to show:

- *CID*—The specified Client ID (CID).
- *ipv6-prefix*—The specified IPv6 prefix.
- *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 dhcpv6 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) Interface VLAN ID or S-VLAN ID interface on which to show binding state information.

interfaces-wildcard—(Optional) 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) Display information about active client bindings for DHCPv6 clients on the specified logical system.

routing-instance *routing-instance-name*—(Optional) Display information about active client bindings for DHCPv6 clients on the specified routing instance.

Required Privilege Level
view

RELATED DOCUMENTATION

<i>Viewing and Clearing DHCP Bindings</i>
clear dhcpv6 server binding 527

List of Sample Output

- [show dhcpv6 server binding on page 595](#)
- [show dhcpv6 server binding detail on page 596](#)
- [show dhcpv6 server binding interface on page 597](#)
- [show dhcpv6 server binding interface detail on page 597](#)
- [show dhcpv6 server binding \(IPv6 Prefix\) on page 598](#)
- [show dhcpv6 server binding \(Session ID\) on page 598](#)
- [show dhcpv6 server binding \(Interfaces VLAN\) on page 598](#)
- [show dhcpv6 server binding \(Interfaces Wildcard\) on page 599](#)
- [show dhcpv6 server binding \(Interfaces Wildcard\) on page 599](#)
- [show dhcpv6 server binding summary on page 599](#)

Output Fields

[Table 17 on page 593](#) lists the output fields for the **show dhcpv6 server binding** command. Output fields are listed in the approximate order in which they appear.

Table 17: show dhcpv6 server binding Output Fields

Field Name	Field Description	Level of Output
<i>number</i> clients, (<i>number</i> init, <i>number</i> bound, <i>number</i> selecting, <i>number</i> requesting, <i>number</i> renewing, <i>number</i> releasing)	Summary counts of the total number of DHCPv6 clients and the number of DHCPv6 clients in each state.	summary
Prefix	Client's DHCPv6 prefix, or prefix used to support multiple address assignment.	brief detail
Session Id	Session ID of the subscriber session.	brief detail

Table 17: show dhcpv6 server binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
Expires	Number of seconds in which lease expires.	brief detail
State	<p>State of the address binding table on the extended DHCPv6 local server:</p> <ul style="list-style-type: none"> • BOUND—Client has active IP address lease. • INIT—Initial state. • RECONFIGURE—Server has sent reconfigure message to client. • RELEASE—Client is releasing IP address lease. • RENEWING—Client sending request to renew IP address lease. • REQUESTING—Client requesting a DHCPv6 server. • SELECTING—Client receiving offers from DHCPv6 servers. 	brief detail
Interface	Interface on which the DHCPv6 request was received.	brief
Client IPv6 Address	Client's IPv6 address.	detail
Client IPv6 Prefix	Client's IPv6 prefix.	detail
Client IPv6 Excluded Prefix	IPv6 Prefix of the DHCP client excluded.	detail
Client DUID	Client's DHCP Unique Identifier (DUID).	brief detail
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
Preferred Lease Expires	Date and UTC time at which the client's IPv6 prefix expires.	detail
Preferred Lease Expires in	Number of seconds at which client's IPv6 prefix expires.	detail
Lease Start	Date and time at which the client's address lease was obtained.	detail
Lease time violated	Lease time violation has occurred.	detail
Incoming Client Interface	Client's incoming interface.	detail

Table 17: show dhcpv6 server binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
Server IP Address	IP address of DHCPv6 server.	detail
Server Interface	Interface of DHCPv6 server.	detail
Client Pool Name	Address pool used to assign IPv6 address.	detail
Client Prefix Pool Name	Address pool used to assign IPv6 prefix.	detail
Client Id length	Length of the DHCPv6 client ID, in bytes.	detail
Client Id	ID of the DHCPv6 client.	detail
Server Id	DHCP unique identifier (DUID) for the DHCPv6 server.	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
Dual Stack Group	DHCPv6 server profile name.	detail
Dual Stack Peer Address	DHCPv6 Peer IP address.	detail

Sample Output

show dhcpv6 server binding

user@host> **show dhcpv6 server binding**

```

Prefix          Session Id  Expires  State  Interface  Client DUID
2001:db8:1111:2222::/64 6          86321    BOUND  ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:01
2001:db8:1111:2222::/64 7          86321    BOUND  ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:02
2001:db8:1111:2222::/64 8          86321    BOUND  ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:03

```

```

2001:db8:1111:2222::/64 9          86321    BOUND    ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:04
2001:db8:1111:2222::/64 10         86321    BOUND    ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:05
2001:db8:2002::1/74 11             86321    BOUND    ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:06

```

show dhcpv6 server binding detail

user@host> show dhcpv6 server binding detail

```

Session Id: 2
  Client IPv6 Prefix:          2001:db8:ffff:0:4::/64
  Client IPv6 Address:        2001:db8:0:8003::1/128
  Client DUID:                LL0x1-00:00:64:01:01:02
  State:                      BOUND(DHCPV6_LOCAL_SERVER_STATE_BOUND)

  Lease Expires:              2016-11-07 08:30:39 PST
  Lease Expires in:           43706 seconds
  Preferred Lease Expires:    2016-11-07 08:30:39 PST
  Preferred Lease Expires in: 43706 seconds
  Lease Start:                2016-11-04 11:00:37 PDT
  Last Packet Received:       2016-11-06 09:00:39 PST
  Incoming Client Interface:  ae0.3221225472
  Client Interface Svlan Id:  2000
  Client Interface Vlan Id:    1
  Server Ip Address:          2001:db8::2
  Server Interface:           none
  Client Profile Name:        my-dual-stack
  Client Id Length:           10
  Client Id:                  /0x00030001/0x00006401/0x0102
  Dual Stack Group:           my-dual-stack
  Dual Stack Peer Address:     192.0.2.10

```

When DHCPv6 binding is configured with prefix exclude option, we get the following output:

user@host> show dhcpv6 server binding detail

```

Session Id: 5
  Client IPv6 Address:        2001:db8:2:3::d/128
  Lease Expires:              2017-12-11 07:45:15 IST
  Lease Expires in:           9999995 seconds

```



```

Preferred Lease Expires:          2017-12-11 07:45:15 IST
Preferred Lease Expires in:       9999995 seconds
Client IPv6 Prefix:              2001:db8::1000:0:0/68
Client IPv6 Excluded Prefix:     2001:db8::1fff:ffff:ff00/120
Lease Expires:                   2017-12-11 07:45:15 IST
Lease Expires in:                9999995 seconds
Preferred Lease Expires:         2017-12-11 07:45:15 IST
Preferred Lease Expires in:      9999995 seconds
Client DUID:                     LL_TIME0x1-0x599553b0-00:10:94:00:00:01

State:                           BOUND(DHCPV6_LOCAL_SERVER_STATE_BOUND)

Lease Start:                     2017-08-17 13:58:32 IST
Last Packet Received:            2017-08-17 13:58:36 IST
Incoming Client Interface:       ge-0/0/0.0
Client Interface Vlan Id:        100
Client Pool Name:                ia_na_pool
Client Prefix Pool Name:         prefix_delegate_pool
Client Id Length:                14
Client Id:                       /0x00010001/0x599553b0/0x00109400/0x0001

Relay Id Length:                 31
Relay Id:
/0x00020000/0x05830130/0x303a3035/0x3a38363a
Relay Id:
/0x34343a65/0x323a6330/0x00000000/0x000000

```

show dhcpv6 server binding interface

```
user@host> show dhcpv6 server binding interface ge-1/0/0:10-101
```

```

Prefix          Session Id  Expires  State   Interface  Client DUID
2001:db8:1111:2222::/64 1      86055    BOUND   ge-1/0/0.100
LL_TIME0x1-0x4b0a53b9-00:10:94:00:00:01

```

show dhcpv6 server binding interface detail

```
user@host> show dhcpv6 server binding interface ge-1/0/0:10-101 detail
```

```

Session Id: 7
Client IPv6 Prefix:          2001:db8:1111:2222::/64
Client DUID:                 LL_TIME0x1-0x2e159c0-00:10:94:00:00:02

```

```

State:                                BOUND(bound)
Lease Expires:                        2009-07-21 10:41:15 PDT
Lease Expires in:                     86136 seconds
Preferred Lease Expires:              2012-07-24 00:18:14 UTC
Preferred Lease Expires in:           600 seconds
Lease Start:                          2009-07-20 10:41:15 PDT
Incoming Client Interface:            ge-1/0/0.0
Server Ip Address:                    0.0.0.0
Server Interface:                     none
Client Id Length:                     14
Client Id:                            /0x00010001/0x02e159c0/0x00109400/0x0002

```

show dhcpv6 server binding (IPv6 Prefix)

user@host> show dhcpv6 server binding 14/0x00010001/0x02b3be8f/0x00109400/0x0005 detail

```

Session Id: 7
  Client IPv6 Prefix:                2001:db8:1111:2222::/64
  Client DUID:                       LL_TIME0x1-0x2e159c0-00:10:94:00:00:02

State:                                BOUND(bound)
Lease Expires:                        2009-07-21 10:41:15 PDT
Lease Expires in:                     86136 seconds
Preferred Lease Expires:              2012-07-24 00:18:14 UTC
Preferred Lease Expires in:           600 seconds
Lease Start:                          2009-07-20 10:41:15 PDT
Incoming Client Interface:            ge-1/0/0.0
Server Ip Address:                    0.0.0.0
Server Interface:                     none
Client Id Length:                     14
Client Id:                            /0x00010001/0x02e159c0/0x00109400/0x0002

```

show dhcpv6 server binding (Session ID)

user@host> show dhcpv6 server binding 8

Prefix	Session Id	Expires	State	Interface	Client DUID
2001:db8::/32	8	86235	BOUND	ge-1/0/0.0	LL_TIME0x1-0x2e159c0-00:10:94:00:00:03

show dhcpv6 server binding (Interfaces VLAN)

user@host> show dhcpv6 server binding ge-1/0/0:100-200

Prefix	Session Id	Expires	State	Interface	Client DUID
2001:db8::/32	11	87583	BOUND	ge-1/0/0.1073741827	
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01					
2001:db8:19::/32	12	87583	BOUND	ge-1/0/0.1073741827	
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01					

show dhcpv6 server binding (Interfaces Wildcard)

```
user@host> show dhcpv6 server binding demux0
```

Prefix	Session Id	Expires	State	Interface	Client DUID
2001:db8::/32	30	79681	BOUND	demux0.1073741824	
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01					
2001:db8:19::/32	31	79681	BOUND	demux0.1073741825	
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01					
2001:db8:C9::/32	32	79681	BOUND	demux0.1073741826	
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01					

show dhcpv6 server binding (Interfaces Wildcard)

```
user@host> show dhcpv6 server binding ge-1/3/*
```

Prefix	Session Id	Expires	State	Interface	Client DUID
2001:db8::/32	22	79681	BOUND	ge-1/3/0.110	
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01					
2001:db8:19::/32	33	79681	BOUND	ge-1/3/0.110	
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01					
2001:db8:C9::/32	24	79681	BOUND	ge-1/3/0.110	
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01					

show dhcpv6 server binding summary

```
user@host> show dhcpv6 server binding summary
```

```
5 clients, (0 init, 5 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)
```

show dhcpv6 server statistics

Syntax

```
show dhcpv6 server statistics
<bulk-leasequery-connections>
<logical-system logical-system-name>
<routing-instance routing-instance-name>
```

Release Information

Command introduced in Junos OS Release 9.6.

bulk-leasequery-connections option introduced in Junos OS Release 16.1.

Description

Display extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server statistics.

Options

bulk-leasequery-connections—(Optional) Display information about DHCPv6 local server bulk leasequery statistics.

logical-system *logical-system-name*—(Optional) Display information about extended DHCPv6 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 DHCPv6 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 dhcpv6 server statistics](#) | 530

DHCPv6 Client MAC Address Validation to Prevent Session Hijacking

List of Sample Output

[show dhcpv6 server statistics on page 602](#)

[show dhcpv6 server statistics bulk-leasequery-connections on page 603](#)

Output Fields

[Table 18 on page 601](#) lists the output fields for the **show dhcpv6 server statistics** command. Output fields are listed in the approximate order in which they appear.

Table 18: show dhcpv6 server statistics Output Fields

Field Name	Field Description
Packets dropped	<p>Number of packets discarded by the extended DHCPv6 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.</p> <ul style="list-style-type: none"> • Total—Total number of packets discarded by the extended DHCPv6 local server • Strict Reconfigure—Number of solicit messages discarded because the client does not support reconfiguration • 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 • Client MAC validation—Number of packets discarded because validation of the client MAC address failed. • Invalid server address—Number of packets discarded because an invalid server address was specified • Lease Time Violation—Number of packets discarded because of a lease time violation • 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 DHCPv6 local server could not send
Advertise Delay	<p>Number of DHCP advertise messages delayed.</p> <ul style="list-style-type: none"> • DELAYED—Number of DHCPv6 advertise packets that have been sent after being delayed. • INPROGRESS—Number of DHCPv6 advertise packets that are in the delay queue. • TOTAL—Total number of delayed DHCPv6 advertise messages; sum of DELAYED and INPROGRESS.

Table 18: show dhcpv6 server statistics Output Fields (*continued*)

Field Name	Field Description
Messages received	<p>Number of DHCPv6 messages received.</p> <ul style="list-style-type: none"> • DHCPV6_CONFIRM—Number of DHCPv6 CONFIRM PDUs received. • DHCPV6_DECLINE—Number of DHCPv6 DECLINE PDUs received. • DHCPV6_INFORMATION_REQUEST—Number of DHCPv6 INFORMATION-REQUEST PDUs received. • DHCPV6_REBIND—Number of DHCPv6 REBIND PDUs received. • DHCPV6_RELAY_FORW—Number of DHCPv6 RELAY-FORW PDUs received. • DHCPV6_RELAY_REPL—Number of DHCPv6 RELAY-REPL PDUs received. • DHCPV6_RELEASE—Number of DHCPv6 RELEASE PDUs received. • DHCPV6_RENEW—Number of DHCPv6 RENEW PDUs received. • DHCPV6_REQUEST—Number of DHCPv6 REQUEST PDUs received. • DHCPV6_SOLICIT—Number of DHCPv6 SOLICIT PDUs received. • DHCPV6_LEASEQUERY—Number of DHCPv6 leasequery messages received.
Messages sent	<p>Number of DHCPv6 messages sent.</p> <ul style="list-style-type: none"> • DHCPV6_ADVERTISE—Number of DHCPv6 ADVERTISE PDUs transmitted. • DHCPV6_REPLY—Number of DHCPv6 ADVERTISE PDUs transmitted. • DHCPV6_LOGICAL_NAK—Number of logical NAK messages sent, signifying T1 and T2 timers with values of zero; subset of DHCPV6_REPLY counter. (Displays only at verbose level.) • DHC6_RECONFIGURE—Number of DHCPv6 RECONFIGURE PDUs transmitted. • DHCPV6_RELAY_REPL—Number of DHCPv6 RELAY-REPL PDUs transmitted. • DHCPV6_RELAY_FORW—Number of DHCPv6 RELAY-FORW PDUs transmitted. • DHCPV6_LEASEQUERY_REPLY—Number of DHCPv6 leasequery replies transmitted to the DHCPv6 relay agent. • DHCPV6_LEASEQUERY_DATA—Number of DHCPv6 LEASEQUERY-DATA packets transmitted. • DHCPV6_LEASEQUERY_DONE—Number of DHCPv6 LEASEQUERY-DONE packets sent.

Sample Output

```
show dhcpv6 server statistics
```

```
user@host> show dhcpv6 server statistics
```

```
Dhcpv6 Packets dropped:
  Total                2
  Lease Time Violation 1
  Client MAC validation 1
```

```
Advertise Delay:
  DELAYED              3
  INPROGRESS           9
  TOTAL                12
```

```
Messages received:
  DHCPV6_DECLINE        0
  DHCPV6_SOLICIT        9
  DHCPV6_INFORMATION_REQUEST 0
  DHCPV6_RELEASE        0
  DHCPV6_REQUEST        5
  DHCPV6_CONFIRM        0
  DHCPV6_RENEW          0
  DHCPV6_REBIND         0
  DHCPV6_RELAY_FORW     0

  DHCPV6_LEASEQUERY     0
```

```
Messages sent:
  DHCPV6_ADVERTISE      9
  DHCPV6_REPLY          5
  DHCPV6_RECONFIGURE    0
  DHCPV6_RELAY_REPL     0
  DHCPV6_LEASEQUERY_REPLY 0
  DHCPV6_LEASEQUERY_DATA 0
  DHCPV6_LEASEQUERY_DONE 0
```

show dhcpv6 server statistics bulk-leasequery-connections

```
user@host> show dhcpv6 server statistics bulk-leasequery-connections
```

```
Total Accepted Connections:          0
Total Not-Accepted Connections:      0
Connections Closed due to Errors:     0
Connections Closed due to max-empty-replies: 0
In-flight Connections:               0
```

show extensible-subscriber-services counters

Syntax

```
show extensible-subscriber-services counters
```

Release Information

Command introduced in Junos OS Release 15.1.

Description

Display number of times various events were processed by Extensible Subscriber Services Manager.

Required Privilege Level

view

RELATED DOCUMENTATION

| *clear extensible-subscriber-services counters*

Output Fields

[Table 19 on page 604](#) describes the output fields for the **show extensible-subscriber-services counters** command. Output fields are listed in the approximate order in which they appear.

Table 19: show extensible-subscriber-services counters Output Fields

Field Name	Field Description
Total Sessions	Number of extensible-subscriber-service control sessions that are currently in the system
Total Services	Number of extensible-subscriber-services that are currently in the system
Active Services	Number of extensible-subscriber-services that are currently in active state
Op-script execution	Total number of operational scripts executed: <ul style="list-style-type: none"> • Successful—Operational scripts that are executed without error • Unsuccessful—Operational scripts that encountered errors
Commit execution	Total number of commit scripts executed: <ul style="list-style-type: none"> • Successful—Commit scripts that are executed without error • Unsuccessful—Commit scripts that encountered errors

Table 19: show extensible-subscriber-services counters Output Fields (*continued*)

Field Name	Field Description
Application execution	Number of applications executed: <ul style="list-style-type: none"> • Successful—Applications that are executed without error • Unsuccessful—Applications that encountered errors
Service requests	Number of service requests received by the daemon. It has counters for total number of requests received, acknowledged (Aked), and negatively acknowledged (Nacked).

Sample Output

```
# show extensible-subscriber-services counters

Total Sessions:      7
Total Services:      14
Active Services:     14

Op-script execution:
  Successful          30
  Unsuccessful        0
Commit execution:
  Successful           4
  Unsuccessful        0
Application execution:
  Successful           0
  Unsuccessful        0
Service requests:
  Received             30
  Aked                 30
  Nacked               0
```

show extensible-subscriber-services debug-information

Syntax

```
show extensible-subscriber-services debug-information
```

Release Information

Command introduced in Junos OS Release 15.1.

Description

Write service session information to a file.

Required Privilege Level

view

Output Fields

[Table 20 on page 606](#) lists the output fields for the **show extensible-subscriber-services debug-information** command.

Table 20: show extensible-subscriber-services debug-information Output Fields

Field Name	Field Description
File Name	Name of the file that contains the service session information

Sample Output

```
#show extensible-subscriber-services debug-information
File Name: /var/tmp/dump_session_record
```

show extensible-subscriber-services dictionary

Syntax

```
show extensible-subscriber-services dictionary
```

Release Information

Command introduced in Junos OS Release 15.1.

Description

Display contents of the dictionary including attributes and services.

Required Privilege Level

view

RELATED DOCUMENTATION

<i>dictionary</i>
<i>request services extensible-subscriber-services reload-dictionary</i>
show extensible-subscriber-services dictionary attributes 612
show extensible-subscriber-services dictionary services 615
Understanding the Dictionary File 88

Output Fields

[Table 21 on page 607](#) lists the output fields for the **show extensible-subscriber-services dictionary** command. Output fields are listed in the approximate order in which they appear.

Table 21: show extensible-subscriber-services dictionary Output Fields

Field Name	Field Description
Acct-Session-Id	Unique accounting identifier that makes it easy to match start and stop records in a log file. The identifier can be in one of the following formats: <ul style="list-style-type: none"> • decimal—For example, 435264 • description—In the generic format, jnpr interface-specifier:subscriber-session-id; For example, jnpr fastEthernet 3/2.6:1010101010101
ERX-Service-Activate	Service to be activated for the subscriber.
Service-Type	Type of service the user has requested or the type of service to be provided.

Table 21: show extensible-subscriber-services dictionary Output Fields (*continued*)

Field Name	Field Description
ADSL-Agent-Remote-Id	Identifier for the subscriber on the digital subscriber line access multiplexer (DSLAM) interface that initiated the service request.
ERX-Med-Port-Number	UDP port in the mediation device to which mirrored traffic is forwarded
NAS-Port-Id	Text string that identifies the physical interface of the NAS that is authenticating the user.
ERX-Med-Dev-Handle	Identifier that associates mirrored traffic with a specific subscriber
ERX-Service-Deactivate	Service to be deactivated for the subscriber.
NAS-IP-Address	IP address of the network access server (NAS) that is requesting authentication of the user.
NAS-Identifier	Name of the NAS that originated the authentication or accounting request.
ERX-Service-Acct-Interval	Amount of time between accounting updates for the service type
ADSL-Agent-Circuit-Id	Identifier for subscriber's access node and the digital subscriber line (DSL) on the access node.
ERX-LI-Action	<p>Traffic mirroring action.</p> <p>For dynamic CoA, VSA 26-58 changes the action on the mirrored traffic identified by VSA 26-59.</p> <p>CoA-Request messages that include any of the RADIUS-based mirroring attributes (VSAs 26-58, 26-59, 26-60, or 26-61) must always include all four VSAs.</p> <p>If the CoA action is to stop mirroring (VSA 26-58 value is 0), then the values of the other three attributes in the CoA message must match the existing attribute values, or the action fails.</p>
ERX-Med-Ip-Address	IP address of content destination device to which mirrored traffic is forwarded.

Sample Output

```
root@ce-bras-mx240-g> show extensible-subscriber-services dictionary
```

RADIUS DICTIONARY

Attribute Name Tag Config Type	Vendor ID	Attribute Code	Attribute Type	Has
Class INVALID	0	25	string	FALSE
Acct-Session-Id INVALID	0	44	string	FALSE
ERX-Service-Activate SERVICE_ACTIVATE	4874	65	string	TRUE
Service-Type INVALID	4874	173	integer	TRUE
ADSL-Agent-Remote-Id INVALID	3561	2	string	FALSE
ERX-Med-Port-Number INVALID	4874	61	OctetString	FALSE
NAS-Port-Id INVALID	0	87	text	FALSE
ERX-Med-Dev-Handle INVALID	4874	59	OctetString	FALSE
ERX-Service-Deactivate SERVICE_DEACTIVATE	4874	66	string	TRUE
NAS-IP-Address INVALID	0	4	address	FALSE
NAS-Identifier INVALID	0	32	string	FALSE
ERX-Service-Acct-Interval INVALID	4874	140	integer	TRUE
ADSL-Agent-Circuit-Id INVALID	3561	1	string	FALSE
ERX-LI-Action SERVICE_ACTIVATE	4874	58	OctetString	FALSE
ERX-Med-Ip-Address INVALID	4874	60	OctetString	FALSE

Services List

Service Name: ngcoco

Service Attribute Name : ERX-Service-Activate

PROVISION ACTION

Action Type : OP-SCRIPT
 Action Version : 1
 Action Name : iceaaa_ngcoco_add

PARAMETERS

Acct-Session-Id
 NAS-Port-Id
 ERX-Service-Activate

DE-PROVISION ACTION

Action Type : OP-SCRIPT
 Action Version : 1
 Action Name : iceaaa_del

PARAMETERS

Acct-Session-Id
 NAS-Port-Id
 ERX-Service-Activate

Service Name: dhcprelay

Service Attribute Name : ERX-Service-Activate

PROVISION ACTION

Action Type : OP-SCRIPT
 Action Version : 1
 Action Name : iceaaa_dhcprelay_add

PARAMETERS

Acct-Session-Id
 NAS-Port-Id
 ERX-Service-Activate

DE-PROVISION ACTION

Action Type : OP-SCRIPT
 Action Version : 1
 Action Name : iceaaa_del

PARAMETERS

Acct-Session-Id
 NAS-Port-Id
 ERX-Service-Activate

Service Name: default

Service Attribute Name : ERX-LI-Action

PROVISION ACTION

Action Type : APPLICATION
 Action Version : 1
 Action Name : LI

PARAMETERS

ERX-Med-Dev-Handle
 ERX-Med-Ip-Address
 ERX-Med-Port-Number

Service Name: dhcprelay

Service Attribute Name : ERX-Service-Deactivate

DE-PROVISION ACTION

Action Type : OP-SCRIPT
 Action Version : 1
 Action Name : iceaaa_del

PARAMETERS

Acct-Session-Id
 NAS-Port-Id
 ERX-Service-Deactivate

Service Name: ngcoco

Service Attribute Name : ERX-Service-Deactivate

DE-PROVISION ACTION

Action Type : OP-SCRIPT
 Action Version : 1
 Action Name : iceaaa_del

PARAMETERS

Acct-Session-Id
 NAS-Port-Id
 ERX-Service-Deactivate

show extensible-subscriber-services dictionary attributes

Syntax

```
show extensible-subscriber-services dictionary attributes
```

Release Information

Command introduced in Junos OS Release 15.1.

Description

Display contents of the dictionary file. It shows only the parameter attributes in the dictionary.

Required Privilege Level

view

RELATED DOCUMENTATION

<i>dictionary</i>
<i>request services extensible-subscriber-services reload-dictionary</i>
show extensible-subscriber-services dictionary 607
show extensible-subscriber-services dictionary services 615
Understanding the Dictionary File 88

Output Fields

[Table 22 on page 612](#) lists the output fields for the **show extensible-subscriber-services dictionary attributes** command. Output fields are listed in the approximate order in which they appear.

Table 22: show extensible-subscriber-services dictionary attributes Output Fields

Field Name	Field Description
Acct-Session-Id	Unique accounting identifier that makes it easy to match start and stop records in a log file. The identifier can be in one of the following formats: <ul style="list-style-type: none"> • decimal—For example, 435264 • description—In the generic format, jnpr interface-specifier:subscriber-session-id; For example, jnpr fastEthernet 3/2.6:1010101010101
ERX-Service-Activate	Service to be activated for the subscriber.
Service-Type	Type of service the user has requested or the type of service to be provided.

Table 22: show extensible-subscriber-services dictionary attributes Output Fields (*continued*)

Field Name	Field Description
ADSL-Agent-Remote-Id	Identifier for the subscriber on the digital subscriber line access multiplexer (DSLAM) interface that initiated the service request.
ERX-Med-Port-Number	UDP port in the mediation device to which mirrored traffic is forwarded
NAS-Port-Id	Text string that identifies the physical interface of the NAS that is authenticating the user.
ERX-Med-Dev-Handle	Identifier that associates mirrored traffic with a specific subscriber
ERX-Service-Deactivate	Service to be deactivated for the subscriber.
NAS-IP-Address	IP address of the network access server (NAS) that is requesting authentication of the user.
NAS-Identifier	Name of the NAS that originated the authentication or accounting request.
ERX-Service-Acct-Interval	Amount of time between accounting updates for the service type
ADSL-Agent-Circuit-Id	Identifier for subscriber's access node and the digital subscriber line (DSL) on the access node.
ERX-LI-Action	<p>Traffic mirroring action.</p> <p>For dynamic CoA, VSA 26-58 changes the action on the mirrored traffic identified by VSA 26-59.</p> <p>CoA-Request messages that include any of the RADIUS-based mirroring attributes (VSAs 26-58, 26-59, 26-60, or 26-61) must always include all four VSAs.</p> <p>If the CoA action is to stop mirroring (VSA 26-58 value is 0), then the values of the other three attributes in the CoA message must match the existing attribute values, or the action fails.</p>
ERX-Med-Ip-Address	IP address of content destination device to which mirrored traffic is forwarded.

Sample Output

```
root@ce-bras-mx240-g> show extensible-subscriber-services dictionary attributes
```

RADIUS DICTIONARY

Attribute Name	Vendor ID	Attribute Code	Attribute Type	Has Tag
Config Type				
Class	0	25	string	FALSE
INVALID				
Acct-Session-Id	0	44	string	FALSE
INVALID				
ERX-Service-Activate	4874	65	string	TRUE
SERVICE_ACTIVATE				
Service-Type	4874	173	integer	TRUE
INVALID				
ADSL-Agent-Remote-Id	3561	2	string	FALSE
INVALID				
ERX-Med-Port-Number	4874	61	OctetString	FALSE
INVALID				
NAS-Port-Id	0	87	text	FALSE
INVALID				
ERX-Med-Dev-Handle	4874	59	OctetString	FALSE
INVALID				
ERX-Service-Deactivate	4874	66	string	TRUE
SERVICE_DEACTIVATE				
NAS-IP-Address	0	4	address	FALSE
INVALID				
NAS-Identifier	0	32	string	FALSE
INVALID				
ERX-Service-Acct-Interval	4874	140	integer	TRUE
INVALID				
ADSL-Agent-Circuit-Id	3561	1	string	FALSE
INVALID				
ERX-LI-Action	4874	58	OctetString	FALSE
SERVICE_ACTIVATE				
ERX-Med-Ip-Address	4874	60	OctetString	FALSE
INVALID				

show extensible-subscriber-services dictionary services

Syntax

```
show extensible-subscriber-services dictionary services
```

Release Information

Command introduced in Junos OS Release 15.1.

Description

Display services configured in the dictionary file and their attributes.

Required Privilege Level

view

RELATED DOCUMENTATION

<i>dictionary</i>
<i>request services extensible-subscriber-services reload-dictionary</i>
show extensible-subscriber-services dictionary 607
show extensible-subscriber-services dictionary attributes 612
Understanding the Dictionary File 88

Output Fields

[Table 23 on page 615](#) lists the output fields for the **show extensible-subscriber-services dictionary services** command. Output fields are listed in the approximate order in which they appear.

Table 23: show extensible-subscriber-services dictionary services Output Fields

Field Name	Field Description
Service Name	Name of the service specified in the dictionary.
Service Attribute Name	Name of the service attribute.
PROVISION ACTION	Name, type, and version of the provisioning action
DE-PROVISION ACTION	Name, type, and version of the deprovisioning action
PARAMETERS	Parameters to be processed for the provisioning or deprovisioning action.

Sample Output

```
root@ce-bras-mx240-g> show extensible-subscriber-services dictionary services
```

RADIUS DICTIONARY

Services List

Service Name: ngcoco

Service Attribute Name : ERX-Service-Activate

PROVISION ACTION

Action Type : OP-SCRIPT

Action Version : 1

Action Name : iceaaa_ngcoco_add

PARAMETERS

Acct-Session-Id

NAS-Port-Id

ERX-Service-Activate

DE-PROVISION ACTION

Action Type : OP-SCRIPT

Action Version : 1

Action Name : iceaaa_del

PARAMETERS

Acct-Session-Id

NAS-Port-Id

ERX-Service-Activate

Service Name: dhcprelay

Service Attribute Name : ERX-Service-Activate

PROVISION ACTION

Action Type : OP-SCRIPT

Action Version : 1

Action Name : iceaaa_dhcprelay_add

PARAMETERS

Acct-Session-Id

NAS-Port-Id

ERX-Service-Activate

DE-PROVISION ACTION

Action Type : OP-SCRIPT
 Action Version : 1
 Action Name : iceaaa_del

PARAMETERS

Acct-Session-Id
 NAS-Port-Id
 ERX-Service-Activate

Service Name: default

Service Attribute Name : ERX-LI-Action

PROVISION ACTION

Action Type : APPLICATION
 Action Version : 1
 Action Name : LI

PARAMETERS

ERX-Med-Dev-Handle
 ERX-Med-Ip-Address
 ERX-Med-Port-Number

Service Name: dhcprelay

Service Attribute Name : ERX-Service-Deactivate

DE-PROVISION ACTION

Action Type : OP-SCRIPT
 Action Version : 1
 Action Name : iceaaa_del

PARAMETERS

Acct-Session-Id
 NAS-Port-Id
 ERX-Service-Deactivate

Service Name: ngcoco

Service Attribute Name : ERX-Service-Deactivate

DE-PROVISION ACTION

Action Type : OP-SCRIPT
 Action Version : 1
 Action Name : iceaaa_del

PARAMETERS

Acct-Session-Id

NAS-Port-Id

ERX-Service-Deactivate

show extensible-subscriber-services sessions

Syntax

```
show extensible-subscriber-services sessions accounting-session-id
```

Release Information

Command introduced in Junos OS Release 15.1.

Description

Display session information. It displays a list of the services applied and their current state. Specify the accounting session ID to view information about a specific session. If an accounting session ID is not specified in the command, the command displays details of all the sessions.

Options

accounting-session-id—Identifier of the session you want information about.

Required Privilege Level

view

RELATED DOCUMENTATION

| *clear extensible-subscriber-services sessions*

Output Fields

[Table 24 on page 619](#) describes the output fields for the **show extensible-subscriber-services sessions** command. Output fields are listed in the approximate order in which they appear.

Table 24: show extensible-subscriber-services sessions Output Fields

Field Name	Field Description
Session ID	Session ID. Accounting session ID must be enclosed in quotation marks.
Timestamp	Time when the session was started
Service Name	Name of the service
Commit State	Commit state: Init , Queued , or Success .
Service ID	Service ID

Sample Output

```
#show extensible-subscriber-services sessions
Session ID: jnpr demux0.1073762028:46422
    Timestamp: Fri Mar  8 04:58:27 2013
    Service Name: ngcoco
    Service Name: dhcprelay

Total Sessions: 1

# show extensible-subscriber-services sessions "jnpr demux0.1073762028:46422"

Service ID: jnpr demux0.1073762028:46422:46425
    Timestamp: Fri Mar  8 04:58:27 2013
    Service Name: ngcoco
    Commit State: Success

Service ID: jnpr demux0.1073762028:46422:46426
    Timestamp: Fri Mar  8 04:58:27 2013
    Service Name: dhcprelay
    Commit State: Success

Total Services: 2
```


show extensible-subscriber-services service

Syntax

```
show extensible-subscriber-services service
```

Release Information

Command introduced in Junos OS Release 15.1.

Description

Display service information. It displays the service name specified in the dictionary, scripts used for provisioning and deprovisioning, and the number of services. A service can have more than one version. Each version is grouped using an action set. The number of action sets is equal to the number of versions for the service. Number of services in all the versions is displayed in **Total Services**.

Required Privilege Level

view

RELATED DOCUMENTATION

[Understanding the Dictionary File](#) | 88

Output Fields

[Table 25 on page 621](#) lists the output fields for the **show extensible-subscriber-services service** command. Output fields are listed in the approximate order in which they appear.

Table 25: show extensible-subscriber-services service Output Fields

Field Name	Field Description
Service Name	Name of the service
Action	Name and number of services provisioned or deprovisioned.
Total Services	Total number of services provisioned or deprovisioned.

Sample Output

```
#show extensible-subscriber-services services
Service Name: ngcoco
  Action:
    Provision:      iceaaa_ngcoco_add_1
    Deprovision:    iceaaa_del_1
    Services:       1301

  Total Services:      1301

Service Name: dhcprelay
  Action:
    Provision:      iceaaa_dhcprelay_add_1
    Deprovision:    iceaaa_del_1
    Services:       1301

  Total Services:      1301
```

show interfaces (Fast Ethernet)

Syntax

```
show interfaces interface-type
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display status information about the specified Fast Ethernet interface.

Options

interface-type—On M Series and T Series routers, the interface type is ***fe-fpc/pic/port***.

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

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

List of Sample Output

[show interfaces \(Fast Ethernet\) on page 641](#)

[show interfaces brief \(Fast Ethernet\) on page 642](#)

[show interfaces detail \(Fast Ethernet\) on page 642](#)

[show interfaces extensive \(Fast Ethernet\) on page 643](#)

Output Fields

[Table 26 on page 624](#) lists the output fields for the **show interfaces** (Fast Ethernet) command. Output fields are listed in the approximate order in which they appear.

Table 26: 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 <i>Common Output Fields Description</i> .	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
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

Table 26: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: Enabled or Disabled .	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> • Online—Autonegotiation is manually configured as online. • Offline—Autonegotiation is manually configured as offline. 	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .	All levels
Link flags	Information about the link. Possible values are described in the “Links Flags” section under <i>Common Output Fields Description</i> .	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	(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

Table 26: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
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	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. <p>Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes.</p> <p>For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 31 under the show interfaces command.</p>	detail extensive

Table 26: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the ignore-l3-incompletes statement. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • Resource errors—Sum of transmit drops. 	extensive

Table 26: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive

Table 26: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Egress queues	<p>Total number of egress queues supported on the specified interface.</p> <p>NOTE: In DPCs that are not of the enhanced type, such as DPC 40x 1GE R, DPCE 20x 1GE + 2x 10GE R, or DPCE 40x 1GE R, you might notice a discrepancy in the output of the show interfaces command because incoming packets might be counted in the Egress queues section of the output. This problem occurs on non-enhanced DPCs because the egress queue statistics are polled from IMQ (Inbound Message Queuing) block of the I-chip. The IMQ block does not differentiate between ingress and egress WAN traffic; as a result, the combined statistics are displayed in the egress queue counters on the Routing Engine. In a simple VPLS scenario, if there is no MAC entry in DMAC table (by sending unidirectional traffic), traffic is flooded and the input traffic is accounted in IMQ. For bidirectional traffic (MAC entry in DMAC table), if the outgoing interface is on the same I-chip then both ingress and egress statistics are counted in a combined way. If the outgoing interface is on a different I-chip or FPC, then only egress statistics are accounted in IMQ. This behavior is expected with non-enhanced DPCs</p>	detail extensive
Queue counters (Egress)	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	detail extensive
Ingress queues	<p>Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.</p>	extensive
Queue counters (Ingress)	<p>CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	extensive

Table 26: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Active alarms and Active defects	<p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the 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.</p> <ul style="list-style-type: none"> • None—There are no active defects or alarms. • Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. 	detail extensive none
OTN FEC statistics	<p>The forward error correction (FEC) counters provide the following statistics:.</p> <ul style="list-style-type: none"> • Corrected Errors—The count of corrected errors in the last second. • Corrected Error Ratio—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits. 	
PCS statistics	<p>(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device.</p> <ul style="list-style-type: none"> • Bit errors—The number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode. • Errored blocks—The number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode. 	detail extensive

Table 26: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> • Total octets and total packets—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see Table 31 under the show interfaces 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

Table 26: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
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 26: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Filter statistics	<p>Receive and Transmit statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected.</p> <ul style="list-style-type: none"> • Input packet count—Number of packets received from the MAC hardware that the filter processed. • Input packet rejects—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • Input DA rejects—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the 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

Table 26: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
PMA PHY	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • PHY Lock—Phase-locked loop • PHY Light—Loss of optical signal 	extensive
WIS section	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B1—Bit interleaved parity for SONET section overhead • SEF—Severely errored framing • LOL—Loss of light • LOF—Loss of frame • ES-S—Errored seconds (section) • SES-S—Severely errored seconds (section) • SEFS-S—Severely errored framing seconds (section) 	extensive

Table 26: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
WIS line	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. State other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B2—Bit interleaved parity for SONET line overhead • REI-L—Remote error indication (near-end line) • RDI-L—Remote defect indication (near-end line) • AIS-L—Alarm indication signal (near-end line) • BERR-SF—Bit error rate fault (signal failure) • BERR-SD—Bit error rate defect (signal degradation) • ES-L—Errored seconds (near-end line) • SES-L—Severely errored seconds (near-end line) • UAS-L—Unavailable seconds (near-end line) • ES-LFE—Errored seconds (far-end line) • SES-LFE—Severely errored seconds (far-end line) • UAS-LFE—Unavailable seconds (far-end line) 	extensive

Table 26: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
WIS path	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B3—Bit interleaved parity for SONET section overhead • REI-P—Remote error indication • LOP-P—Loss of pointer (path) • AIS-P—Path alarm indication signal • RDI-P—Path remote defect indication • UNEQ-P—Path unequipped • PLM-P—Path payload (signal) label mismatch • ES-P—Errored seconds (near-end STS path) • SES-P—Severely errored seconds (near-end STS path) • UAS-P—Unavailable seconds (near-end STS path) • SES-PFE—Severely errored seconds (far-end STS path) • UAS-PFE—Unavailable seconds (far-end STS path) 	extensive

Table 26: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Autonegotiation information	<p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> • Negotiation status: <ul style="list-style-type: none"> • Incomplete—Ethernet interface has the speed or link mode configured. • No autonegotiation—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation. • Complete—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner status—OK when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner: <ul style="list-style-type: none"> • Link mode—Depending on the capability of the attached Ethernet device, either Full-duplex or Half-duplex. • Flow control—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is None. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports both PAUSE on receive and transmit or only PAUSE receive). • Remote fault—Remote fault information from the link partner—Failure indicates a receive link error. OK indicates that the link partner is receiving. Negotiation error indicates a negotiation error. Offline indicates that the link partner is going offline. • Local resolution—Information from the link partner: <ul style="list-style-type: none"> • Flow control—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports both PAUSE on receive and transmit or only PAUSE receive). • Remote fault—Remote fault information. Link OK (no error detected on receive), Offline (local interface is offline), and Link Failure (link error detected on receive). 	extensive

Table 26: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
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: <ul style="list-style-type: none"> • Destination slot—FPC slot number. 	extensive
CoS information	Information about the CoS queue for the physical interface. <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none

Table 26: 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
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> .	All levels
VLAN-Tag	<p>Rewrite profile applied to incoming or outgoing frames on the outer (Out) VLAN tag or for both the outer and inner (In) VLAN tags.</p> <ul style="list-style-type: none"> • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • pop—The outer VLAN tag of the incoming frame is removed. • swap—The outer VLAN tag of the incoming frame is overwritten with the user specified VLAN tag information. • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • push-push—Two VLAN tags are pushed in from the incoming frame. • swap-push—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • swap-swap—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user specified VLAN tag value. • pop-swap—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • pop-pop—Both the outer and inner VLAN tags of the incoming frame are removed. 	brief detail extensive none
Demux:	<p>IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following:</p> <ul style="list-style-type: none"> • Source Family Inet • Destination Family Inet 	detail extensive none
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under <i>Common Output Fields Description</i> .	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none

Table 26: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	detail extensive none
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface set • 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	<p>Number and rate of bytes and packets transiting the switch.</p> <p>NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the Output bytes and Output packets interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.</p>	extensive
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive none
Flags	Information about protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	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

Table 26: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
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 <i>Common Output Fields Description</i> .	detail extensive none
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
Flags	Information about address flag (possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	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: 00:00:5e:00:53:38, Hardware address: 00:00:5e:00:53:38
Last flapped    : 2006-01-20 14:50:58 PST (2w4d 00:44 ago)
Input rate      : 0 bps (0 pps)
Output rate     : 0 bps (0 pps)
Active alarms   : None
Active defects  : None
Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198)
  Flags: SNMP-Traps Encapsulation: ENET2
  Protocol inet, MTU: 1500
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 203.0.113/24, Local: 203.0.113.1, Broadcast: 203.0.113.255

```

show interfaces brief (Fast Ethernet)

user@host> **show interfaces fe-0/0/0 brief**

```

Physical interface: fe-0/0/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
Device flags      : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Logical interface fe-0/0/0.0
  Flags: SNMP-Traps Encapsulation: ENET2
  inet 203.0.113.1/24

```

show interfaces detail (Fast Ethernet)

user@host> **show interfaces fe-0/0/0 detail**

```

Physical interface: fe-0/0/0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 22, Generation: 5391
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
Device flags      : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
CoS queues       : 4 supported, 4 maximum usable queues
Hold-times       : Up 0 ms, Down 0 ms
Current address: 00:00:5e:00:53:38, Hardware address: 00:00:5e:00:53:3f:38
Last flapped    : 2006-01-20 14:50:58 PST (2w4d 00:45 ago)

```

```

Statistics last cleared: Never
Traffic statistics:
  Input  bytes   :                0          0 bps
  Output bytes   :               42          0 bps
  Input  packets :                0          0 pps
  Output packets :                1          0 pps
Active alarms   : None
Active defects  : None
Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198) (Generation 67)
  Flags: SNMP-Traps Encapsulation: ENET2
  Protocol inet, MTU: 1500, Generation: 105, Route table: 0
    Flags: Is-Primary, Mac-Validate-Strict
    Mac-Validate Failures: Packets: 0, Bytes: 0
    Addresses, Flags: Is-Preferred Is-Primary
    Destination: 203.0.113/24, Local: 203.0.113.1, Broadcast: 203.0.113.255,
      Generation: 136

```

show interfaces extensive (Fast Ethernet)

user@host> **show interfaces fe-0/0/0 extensive**

```

Physical interface: fe-0/0/0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 22, Generation: 5391
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed:
  100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:38, Hardware address: 00:00:5e:00:53:38
  Last flapped   : 2006-01-20 14:50:58 PST (2w4d 00:46 ago)
  Statistics last cleared: Never
Traffic statistics:
  Input  bytes   :                0          0 bps
  Output bytes   :               42          0 bps
  Input  packets :                0          0 pps
  Output packets :                1          0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 3, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

```

```

    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Active alarms   : None
Active defects  : None
MAC statistics:
    Receive      Transmit
Total octets    0          64
Total packets   0          1
Unicast packets 0          0
Broadcast packets 0        1
Multicast packets 0        0
CRC/Align errors 0          0
FIFO errors      0          0
MAC control frames 0          0
MAC pause frames 0          0
Oversized frames 0
Jabber frames    0
Fragment frames  0
VLAN tagged frames 0
Code violations   0
Filter statistics:
    Input packet count      0
    Input packet rejects    0
    Input DA rejects        0
    Input SA rejects        0
    Output packet count     1
    Output packet pad count 0
    Output packet error count 0
    CAM destination filters: 1, CAM source filters: 0
Autonegotiation information:
    Negotiation status: Complete
    Link partner:
        Link partner: Full-duplex, Flow control: None, Remote fault: Ok
    Local resolution:
Packet Forwarding Engine configuration:
    Destination slot: 0
CoS information:
    Bandwidth      Buffer Priority  Limit
                   %      bps    %      usec
0 best-effort     95     950000000  95      0      low  none
3 network-control  5      500000000   5      0      low  none
Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198) (Generation 67)
Flags: SNMP-Traps Encapsulation: ENET2
Protocol inet, MTU: 1500, Generation: 105, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary

```


Destination: 203.0.113/24, Local: 203.0.113.1, Broadcast: 203.0.113.255,
Generation: 136

show interfaces

List of Syntax

[Syntax \(Gigabit Ethernet\) on page 646](#)

[Syntax \(10 Gigabit Ethernet\) on page 646](#)

[Syntax \(SRX Series Devices and \(vSRX and vSRX 3.0 platforms\)\) on page 646](#)

Syntax (Gigabit Ethernet)

```
show interfaces ge-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Syntax (10 Gigabit Ethernet)

```
show interfaces xe-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Syntax (SRX Series Devices and (vSRX and vSRX 3.0 platforms))

```
show interfaces (
  <interface-name>
  <brief | detail | extensive | terse>
  <controller interface-name>|
  <descriptions interface-name>|
  <destination-class (all | destination-class-name logical-interface-name)>|
  <diagnostics optics interface-name>|
  <far-end-interval interface-fpc/pic/port>|
  <filters interface-name>|
  <flow-statistics interface-name>|
  <interval interface-name>|
  <load-balancing (detail | interface-name)>|
  <mac-database mac-address mac-address>|
  <mc-ae id identifier unit number revertive-info>|
  <media interface-name>|
  <policers interface-name>|
```

```

<queue both-ingress-egress egress forwarding-class forwarding-class ingress l2-statistics>|
<redundancy (detail | interface-name)>|
<routing brief detail summary interface-name>|
<routing-instance (all | instance-name)>|
<snmp-index snmp-index>|
<source-class (all | destination-class-name logical-interface-name)>|
<statistics interface-name>|
<switch-port switch-port number>|
<transport pm (all | optics | otn) (all | current | currentday | interval | previousday) (all | interface-name)>|
<zone interface-name>
)

```

Release Information

Command introduced before Junos OS Release 7.4 for Gigabit interfaces.

Command introduced in Junos OS Release 8.0 for 10 Gigabit interfaces.

Command modified in Junos OS Release 9.5 for SRX Series devices.

Command introduced in Junos OS Release 18.1 for Gigabit interfaces.

Command modified in Junos OS Release 19.3R1 for MX Series Routers.

Starting in Junos OS Release 19.3R1, Output fields **lindex** and **speed** is modified in the **show interfaces *interface name* extensive** command, on all MX Series routers.

- The default behavior of WAN-PHY interface remains the same. The new **precise-bandwidth** option reflects the new speed (9.294-Gbps) configured on the supported line cards.
- The WAN-PHY framing mode is supported only on MPC5E and MPC6E line cards.

Starting in Junos OS Release 19.3R1, class of service (CoS) features can be configured on the physical interface with speed rates of 1-Gbps, 10-Gbps, 40-Gbps, and 100-Gbps to provide better bandwidth for processing traffic during congestion using variant speeds.

Description

Display status information about the specified Gigabit Ethernet interface.

(M320, M120, MX Series, and T Series routers only) Display status information about the specified 10-Gigabit Ethernet interface.

Display the IPv6 interface traffic statistics about the specified Gigabit Ethernet interface for MX series routers. The input and output bytes (bps) and packets (pps) rates are not displayed for IFD and local traffic.

Display status information and statistics about interfaces on SRX Series, vSRX, and vSRX 3.0 platforms running Junos OS.

NOTE: On SRX Series appliances, on configuring identical IPs on a single interface, you will not see a warning message; instead, you will see a syslog message.

Starting in Junos OS Release 18.4R1, Output fields **Next-hop** and **vpls-status** is displayed in the **show interfaces *interface name* detail** command, only for Layer 2 protocols on MX480 routers.

Options

For Gigabit interfaces:

ge-*fpc/pic/port*—Display standard information about the specified Gigabit Ethernet interface.

NOTE: Interfaces with different speeds are named uniformly with ge-0/0/x for backward compatibility. Use the **show interfaces** command to view the interface speeds.

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

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

For 10 Gigabit interfaces:

xe-*fpc/pic/port*—Display standard information about the specified 10-Gigabit Ethernet interface.

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

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

For SRX interfaces:

- **interface-name**—(Optional) Display standard information about the specified interface. Following is a list of typical interface names. Replace pim with the PIM slot and port with the port number.

- **at-*pim*/0/*port***—ATM-over-ADSL or ATM-over-SHDSL interface.
- **ce1-*pim*/0/ *port***—Channelized E1 interface.
- **cl-0/0/8**—3G wireless modem interface for SRX320 devices.
- **ct1-*pim*/0/*port***—Channelized T1 interface.
- **dl0**—Dialer Interface for initiating ISDN and USB modem connections.
- **e1-*pim*/0/*port***—E1 interface.
- **e3-*pim*/0/*port***—E3 interface.
- **fe-*pim*/0/*port***—Fast Ethernet interface.
- **ge-*pim*/0/*port***—Gigabit Ethernet interface.
- **se-*pim*/0/*port***—Serial interface.
- **t1-*pim*/0/*port***—T1 (also called DS1) interface.
- **t3-*pim*/0/*port***—T3 (also called DS3) interface.
- **wx-slot/0/0**—WAN acceleration interface, for the WXC Integrated Services Module (ISM 200).
- **interface-name**—(Optional) Display standard information about the specified interface. Following is a list of typical interface names. Replace *pim* with the PIM slot and *port* with the port number.
 - **at-*pim*/0/*port***—ATM-over-ADSL or ATM-over-SHDSL interface.
 - **ce1-*pim*/0/ *port***—Channelized E1 interface.
 - **cl-0/0/8**—3G wireless modem interface for SRX320 devices.
 - **ct1-*pim*/0/*port***—Channelized T1 interface.
 - **dl0**—Dialer Interface for initiating ISDN and USB modem connections.
 - **e1-*pim*/0/*port***—E1 interface.
 - **e3-*pim*/0/*port***—E3 interface.
 - **fe-*pim*/0/*port***—Fast Ethernet interface.
 - **ge-*pim*/0/*port***—Gigabit Ethernet interface.
 - **se-*pim*/0/*port***—Serial interface.
 - **t1-*pim*/0/*port***—T1 (also called DS1) interface.
 - **t3-*pim*/0/*port***—T3 (also called DS3) interface.
 - **wx-slot/0/0**—WAN acceleration interface, for the WXC Integrated Services Module (ISM 200).

Additional Information

In a logical system, this command displays information only about the logical interfaces and not about the physical interfaces.

Required Privilege Level

view

Release History Table

Release	Description
18.4R1	Starting in Junos OS Release 18.4R1, Output fields Next-hop and vpls-status is displayed in the show interfaces <i>interface name</i> detail command, only for Layer 2 protocols on MX480 routers.

RELATED DOCUMENTATION

[Understanding Layer 2 Interfaces on Security Devices](#)

[Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration](#)

[Verifying and Managing Configurations for Dynamic VLANs Based on Access-Line Identifiers](#)
List of Sample Output
[show interfaces \(Gigabit Ethernet\) on page 699](#)
[show interfaces \(Gigabit Ethernet on MX Series Routers\) on page 700](#)
[show interfaces \(link degrade status\) on page 701](#)
[show interfaces extensive \(Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration\) on page 702](#)
[show interfaces brief \(Gigabit Ethernet\) on page 703](#)
[show interfaces detail \(Gigabit Ethernet\) on page 703](#)
[show interfaces extensive \(Gigabit Ethernet IQ2\) on page 705](#)
[show interfaces \(Gigabit Ethernet Unnumbered Interface\) on page 709](#)
[show interfaces \(ACI Interface Set Configured\) on page 709](#)
[show interfaces \(ALI Interface Set\) on page 710](#)
[show interfaces extensive \(10-Gigabit Ethernet, LAN PHY Mode, IQ2\) on page 710](#)
[show interfaces extensive \(10-Gigabit Ethernet, WAN PHY Mode\) on page 713](#)
[show interfaces extensive \(10-Gigabit Ethernet, DWDM OTN PIC\) on page 716](#)
[show interfaces extensive \(10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode\) on page 719](#)
[show interfaces extensive \(10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only\) on page 720](#)
[show interfaces extensive \(10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only\) on page 721](#)
[Sample Output SRX Gigabit Ethernet on page 723](#)
[Sample Output SRX Gigabit Ethernet on page 723](#)
[show interfaces \(Gigabit Ethernet for vSRX and vSRX 3.0\) on page 724](#)
[show interfaces detail \(Gigabit Ethernet\) on page 725](#)
[show interfaces statistics st0.0 detail on page 727](#)

[show interfaces extensive \(Gigabit Ethernet\) on page 728](#)
[show interfaces terse on page 731](#)
[show interfaces terse \(vSRX and vSRX 3.0\) on page 732](#)
[show interfaces controller \(Channelized E1 IQ with Logical E1\) on page 733](#)
[show interfaces controller \(Channelized E1 IQ with Logical DS0\) on page 733](#)
[show interfaces descriptions on page 733](#)
[show interfaces destination-class all on page 734](#)
[show interfaces diagnostics optics on page 734](#)
[show interfaces far-end-interval coc12-5/2/0 on page 735](#)
[show interfaces far-end-interval coc1-5/2/1:1 on page 736](#)
[show interfaces filters on page 736](#)
[show interfaces flow-statistics \(Gigabit Ethernet\) on page 737](#)
[show interfaces interval \(Channelized OC12\) on page 738](#)
[show interfaces interval \(E3\) on page 739](#)
[show interfaces interval \(SONET/SDH\) \(SRX devices\) on page 739](#)
[show interfaces load-balancing \(SRX devices\) on page 740](#)
[show interfaces load-balancing detail \(SRX devices\) on page 740](#)
[show interfaces mac-database \(All MAC Addresses on a Port SRX devices\) on page 740](#)
[show interfaces mac-database \(All MAC Addresses on a Service SRX devices\) on page 741](#)
[show interfaces mac-database mac-address on page 742](#)
[show interfaces mc-ae \(SRX devices\) on page 742](#)
[show interfaces media \(SONET/SDH\) on page 742](#)
[show interfaces policers \(SRX devices\) on page 743](#)
[show interfaces policers interface-name \(SRX devices\) on page 744](#)
[show interfaces queue \(SRX devices\) on page 744](#)
[show interfaces redundancy \(SRX devices\) on page 745](#)
[show interfaces redundancy \(Aggregated Ethernet SRX devices\) on page 745](#)
[show interfaces redundancy detail \(SRX devices\) on page 746](#)
[show interfaces routing brief \(SRX devices\) on page 746](#)
[show interfaces routing detail \(SRX devices\) on page 747](#)
[show interfaces routing-instance all \(SRX devices\) on page 747](#)
[show interfaces snmp-index \(SRX devices\) on page 748](#)
[show interfaces source-class all \(SRX devices\) on page 748](#)
[show interfaces statistics \(Fast Ethernet SRX devices\) on page 749](#)
[show interfaces switch-port \(SRX devices\) on page 750](#)
[show interfaces transport pm \(SRX devices\) on page 750](#)
[show security zones \(SRX devices\) on page 752](#)

Output Fields

[Table 27 on page 652](#) 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 28 on page 690](#).

Table 27: 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 <i>Common Output Fields Description</i> .	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

Table 27: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> • Online—Autonegotiation is manually configured as online. • Offline—Autonegotiation is manually configured as offline. 	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .	All levels
Link flags	Information about the link. Possible values are described in the “Links Flags” section under <i>Common Output Fields Description</i> .	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

Table 27: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output Rate	Output rate in bps and pps. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Egress account overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for egress traffic.	detail extensive
Ingress account overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for ingress traffic.	detail extensive
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Output bytes—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. <p>Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes.</p> <p>For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 31 under the show interfaces command.</p>	detail extensive

Table 27: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the ignore-l3-incompletes statement. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • Resource errors—Sum of transmit drops. 	extensive

Table 27: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. <p>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.</p> <ul style="list-style-type: none"> • Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number must always be 0. If it is nonzero, there is a software bug. • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field must never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive

Table 27: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
Egress queues	<p>Total number of egress queues supported on the specified interface.</p> <p>NOTE: In DPCs that are not of the enhanced type, such as DPC 40x 1GE R, DPCE 20x 1GE + 2x 10GE R, or DPCE 40x 1GE R, you might notice a discrepancy in the output of the show interfaces command because incoming packets might be counted in the Egress queues section of the output. This problem occurs on non-enhanced DPCs because the egress queue statistics are polled from IMQ (Inbound Message Queuing) block of the I-chip. The IMQ block does not differentiate between ingress and egress WAN traffic; as a result, the combined statistics are displayed in the egress queue counters on the Routing Engine. In a simple VPLS scenario, if there is no MAC entry in DMAC table (by sending unidirectional traffic), traffic is flooded and the input traffic is accounted in IMQ. For bidirectional traffic (MAC entry in DMAC table), if the outgoing interface is on the same I-chip then both ingress and egress statistics are counted in a combined way. If the outgoing interface is on a different I-chip or FPC, then only egress statistics are accounted in IMQ. This behavior is expected with non-enhanced DPCs</p>	detail extensive
Queue counters (Egress)	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. <p>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.</p>	detail extensive
Ingress queues	<p>Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.</p>	extensive
Queue counters (Ingress)	<p>CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	extensive

Table 27: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Active alarms and Active defects	<p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value None or Link.</p> <ul style="list-style-type: none"> • None—There are no active defects or alarms. • Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. 	detail extensive none
Interface transmit statistics	<p>(On MX Series devices) Status of the interface-transmit-statistics configuration: Enabled or Disabled.</p> <ul style="list-style-type: none"> • Enabled—When the interface-transmit-statistics statement is included in the configuration. If this is configured, the interface statistics show the actual transmitted load on the interface. • Disabled—When the interface-transmit-statistics statement is not included in the configuration. If this is not configured, the interface statistics show the offered load on the interface. 	detail extensive
OTN FEC statistics	<p>The forward error correction (FEC) counters provide the following statistics:</p> <ul style="list-style-type: none"> • Corrected Errors—Count of corrected errors in the last second. • Corrected Error Ratio—Corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits. 	detail extensive
PCS statistics	<p>(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device.</p> <ul style="list-style-type: none"> • Bit errors—Number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode. • Errorred blocks—Number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode. 	detail extensive

Table 27: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Link Degrad	<p>Shows the link degrade status of the physical link and the estimated bit error rates (BERs). This field is available only for the PICs supporting the physical link monitoring feature.</p> <ul style="list-style-type: none"> • Link Monitoring—Indicates if physical link degrade monitoring is enabled on the interface. <ul style="list-style-type: none"> • Enable—Indicates that link degrade monitoring has been enabled (using the link-degrade-monitor statement) on the interface. • Disable—Indicates that link degrade monitoring has not been enabled on the interface. If link degrade monitoring has not been enabled, the output does not show any related information, such as BER values and thresholds. • Link Degrad Set Threshold—The BER threshold value at which the link is considered degraded and a corrective action is triggered. • Link Degrad Clear Threshold—The BER threshold value at which the degraded link is considered recovered and the corrective action applied to the interface is reverted. • Estimated BER—The estimated bit error rate. • Link-degrade event—Shows link degrade event information. <ul style="list-style-type: none"> • Seconds—Time (in seconds) elapsed after a link degrade event occurred. • Count—The number of link degrade events recorded. • State—Shows the link degrade status (example: Defect Active). 	detail extensive

Table 27: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
MAC statistics		extensive

Table 27: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> • Total octets and total packets—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see Table 31 under the show interfaces command. • Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets. • CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • FIFO error—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • MAC control frames—Number of MAC control frames. • MAC pause frames—Number of MAC control frames with pause operational code. • Oversized frames—There are two possible conditions regarding the number of oversized frames: <ul style="list-style-type: none"> • Packet length exceeds interface MTU, 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. 	

Table 27: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
	<p>NOTE: The 20-port Gigabit Ethernet MIC (MIC-3D-20GE-SFP) does not have hardware counters for VLAN frames. Therefore, the VLAN tagged frames field displays 0 when the show interfaces command is executed on a 20-port Gigabit Ethernet MIC. In other words, the number of VLAN tagged frames cannot be determined for the 20-port Gigabit Ethernet MIC.</p> <ul style="list-style-type: none"> • Code violations—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.” 	
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 27: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Filter statistics	<p>Receive and Transmit statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet may enter the system or be rejected.</p> <ul style="list-style-type: none"> • Input packet count—Number of packets received from the MAC hardware that the filter processed. • Input packet rejects—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • Input DA rejects—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the router from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local router (which the router is rejecting). • Input SA rejects—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field must increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • Output packet count—Number of packets that the filter has given to the MAC hardware. • Output packet pad count—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • Output packet error count—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field must not increment. • CAM destination filters, CAM source filters—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields must be 0. 	extensive

Table 27: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
PMA PHY	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • PHY Lock—Phase-locked loop • PHY Light—Loss of optical signal 	extensive
WIS section	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B1—Bit interleaved parity for SONET section overhead • SEF—Severely errored framing • LOL—Loss of light • LOF—Loss of frame • ES-S—Errored seconds (section) • SES-S—Severely errored seconds (section) • SEFS-S—Severely errored framing seconds (section) 	extensive

Table 27: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
WIS line	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B2—Bit interleaved parity for SONET line overhead • REI-L—Remote error indication (near-end line) • RDI-L—Remote defect indication (near-end line) • AIS-L—Alarm indication signal (near-end line) • BERR-SF—Bit error rate fault (signal failure) • BERR-SD—Bit error rate defect (signal degradation) • ES-L—Errored seconds (near-end line) • SES-L—Severely errored seconds (near-end line) • UAS-L—Unavailable seconds (near-end line) • ES-LFE—Errored seconds (far-end line) • SES-LFE—Severely errored seconds (far-end line) • UAS-LFE—Unavailable seconds (far-end line) 	extensive

Table 27: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
WIS path	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B3—Bit interleaved parity for SONET section overhead • REI-P—Remote error indication • LOP-P—Loss of pointer (path) • AIS-P—Path alarm indication signal • RDI-P—Path remote defect indication • UNEQ-P—Path unequipped • PLM-P—Path payload (signal) label mismatch • ES-P—Errored seconds (near-end STS path) • SES-P—Severely errored seconds (near-end STS path) • UAS-P—Unavailable seconds (near-end STS path) • SES-PFE—Severely errored seconds (far-end STS path) • UAS-PFE—Unavailable seconds (far-end STS path) 	extensive

Table 27: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Autonegotiation information	<p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> • Negotiation status: <ul style="list-style-type: none"> • Incomplete—Ethernet interface has the speed or link mode configured. • No autonegotiation—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation. • Complete—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner status—OK when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner—Information from the remote Ethernet device: <ul style="list-style-type: none"> • Link mode—Depending on the capability of the link partner, either Full-duplex or Half-duplex. • Flow control—Types of flow control supported by the link partner. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), Symmetric/Asymmetric (link partner supports PAUSE on receive and transmit or only PAUSE on transmit), and None (link partner does not support flow control). • Remote fault—Remote fault information from the link partner—Failure indicates a receive link error. OK indicates that the link partner is receiving. Negotiation error indicates a negotiation error. Offline indicates that the link partner is going offline. • Local resolution—Information from the local Ethernet device: <ul style="list-style-type: none"> • Flow control—Types of flow control supported by the local device. For Gigabit Ethernet interfaces, advertised capabilities are Symmetric/Asymmetric (local device supports PAUSE on receive and transmit or only PAUSE on receive) and None (local device does not support flow control). Depending on the result of the negotiation with the link partner, local resolution flow control type will display Symmetric (local device supports PAUSE on receive and transmit), Asymmetric (local device supports PAUSE on receive), and None (local device does not support flow control). • Remote fault—Remote fault information. Link OK (no error detected on receive), Offline (local interface is offline), and Link Failure (link error detected on receive). 	extensive

Table 27: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
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: <ul style="list-style-type: none"> • Destination slot—FPC slot number. 	extensive
CoS information	Information about the CoS queue for the physical interface. <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Table 27: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> .	All levels
VLAN-Tag	<p>Rewrite profile applied to incoming or outgoing frames on the outer (Out) VLAN tag or for both the outer and inner (In) VLAN tags.</p> <ul style="list-style-type: none"> • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • pop—The outer VLAN tag of the incoming frame is removed. • swap—The outer VLAN tag of the incoming frame is overwritten with the user-specified VLAN tag information. • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • push-push—Two VLAN tags are pushed in from the incoming frame. • swap-push—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • swap-swap—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user-specified VLAN tag value. • pop-swap—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • pop-pop—Both the outer and inner VLAN tags of the incoming frame are removed. 	brief detail extensive none
Demux	<p>IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following:</p> <ul style="list-style-type: none"> • Source Family Inet • Destination Family Inet 	detail extensive none
Encapsulation	Encapsulation on the logical interface.	All levels

Table 27: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
ACI VLAN	<p>Information displayed for agent circuit identifier (ACI) interface set configured with the agent-circuit-id autoconfiguration stanza.</p> <p>Dynamic Profile—Name of the dynamic profile that defines the ACI interface set.</p> <p>If configured, the ACI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ACI information.</p> <p>NOTE: The ACI VLAN field is replaced with the Line Identity field when an ALI interface set is configured with the line-identity autoconfiguration stanza.</p>	brief detail extensive none
Line Identity	<p>Information displayed for access-line-identifier (ALI) interface sets configured with the line-identity autoconfiguration stanza.</p> <ul style="list-style-type: none"> • Dynamic Profile—Name of the dynamic profile that defines the ALI interface set. • Trusted option used to create the ALI interface set: Circuit-id, Remote-id, or Accept-no-ids. More than one option can be configured. <p>If configured, the ALI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ALI information.</p> <p>NOTE: The Line Identity field is replaced with the ACI VLAN field when an ACI interface set is configured with the agent-circuit-id autoconfiguration stanza.</p>	detail
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under <i>Common Output Fields Description</i> .	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none
Neighbor Discovery Protocol (NDP) Queue Statistics	<p>NDP statistics for protocol inet6 under logical interface statistics.</p> <ul style="list-style-type: none"> • Max nh cache—Maximum interface neighbor discovery nexthop cache size. • New hold nh limit—Maximum number of new unresolved nexthops. • Curr nh cnt—Current number of resolved nexthops in the NDP queue. • Curr new hold cnt—Current number of unresolved nexthops in the NDP queue. • NH drop cnt—Number of NDP requests not serviced. 	All levels

Table 27: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Dynamic Profile	Name of the dynamic profile that was used to create this interface configured with a Point-to-Point Protocol over Ethernet (PPPoE) family.	detail extensive none
Service Name Table	Name of the service name table for the interface configured with a PPPoE family.	detail extensive none
Max Sessions	Maximum number of PPPoE logical interfaces that can be activated on the underlying interface.	detail extensive none
Duplicate Protection	State of PPPoE duplicate protection: On or Off . When duplicate protection is configured for the underlying interface, a dynamic PPPoE logical interface cannot be activated when an existing active logical interface is present for the same PPPoE client.	detail extensive none
Direct Connect	State of the configuration to ignore DSL Forum VSAs: On or Off . When configured, the router ignores any of these VSAs received from a directly connected CPE device on the interface.	detail extensive none
AC Name	Name of the access concentrator.	detail extensive none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	detail extensive none
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Input packets, Output packets—Number of packets received and transmitted on the interface set. 	detail extensive
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	extensive
Local statistics	Number and rate of bytes and packets destined to the router.	extensive

Table 27: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Transit statistics	<p>Number and rate of bytes and packets transiting the switch.</p> <p>NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the Output bytes and Output packets interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.</p>	extensive
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive none
Flags	Information about protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	detail extensive
Donor interface	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	detail extensive none
Preferred source address	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	detail extensive none
Input Filters	Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parentheses next to all interfaces.	detail extensive
Output Filters	Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parentheses next to all interfaces.	detail extensive
Mac-Validate Failures	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	detail extensive none
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	detail extensive none

Table 27: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
Flags	Information about the address flag. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	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

The following table describes the output fields for the **show interfaces** (10-Gigabit Ethernet) command.

Field Name	Field Description	Level of Output
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Loopback	Loopback status: Enabled or Disabled . If loopback is enabled, type of loopback: Local or Remote .	All levels

Source filtering	Source filtering status: Enabled or Disabled .	All levels
LAN-PHY mode	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
WAN-PHY mode	10-Gigabit Ethernet interface operating in Wide Area Network Physical Layer Device (WAN PHY) mode. WAN PHY allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and other devices intended for SONET/SDH.	All levels
Unidirectional	Unidirectional link mode status for 10-Gigabit Ethernet interface: Enabled or Disabled for parent interface; Rx-only or Tx-only for child interfaces.	All levels
Flow control	Flow control status: Enabled or Disabled .	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: Enabled or Disabled .	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> • Online—Autonegotiation is manually configured as online. • Offline—Autonegotiation is manually configured as offline. 	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .	All levels
Link flags	Information about the link. Possible values are described in the “Links Flags” section under <i>Common Output Fields Description</i> .	All levels
Wavelength	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels
Frequency	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels
CoS queues	Number of CoS queues configured.	detail extensive none
Schedulers	(Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces only) Number of CoS schedulers configured.	extensive
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive

Current address	Configured MAC address.	detail extensive none
Hardware address	Hardware MAC address.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive none
Input Rate	Input rate in bits per second (bps) and packets per second (pps). The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None specified
Output Rate	Output rate in bps and pps. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Egress account overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for egress traffic.	detail extensive
Ingress account overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for ingress traffic.	detail extensive
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Output bytes—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive

Input errors

Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:

- **Errors**—Sum of the incoming frame aborts and FCS errors.
- **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.
- **Framing errors**—Number of packets received with an invalid frame checksum (FCS).
- **Runts**—Number of frames received that are smaller than the runt threshold.
- **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle.
- **L3 incompletes**—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the **ignore-l3-incompletes** statement.
- **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame.
- **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.
- **FIFO errors**—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.
- **Resource errors**—Sum of transmit drops.

extensive

Output errors

Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:

extensive

- **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.

Egress queues	<p>Total number of egress queues supported on the specified interface.</p> <p>NOTE: In DPCs that are not of the enhanced type, such as DPC 40x 1GE R, DPCE 20x 1GE + 2x 10GE R, or DPCE 40x 1GE R, you might notice a discrepancy in the output of the show interfaces command because incoming packets might be counted in the Egress queues section of the output. This problem occurs on non-enhanced DPCs because the egress queue statistics are polled from IMQ (Inbound Message Queuing) block of the I-chip. The IMQ block does not differentiate between ingress and egress WAN traffic; as a result, the combined statistics are displayed in the egress queue counters on the Routing Engine. In a simple VPLS scenario, if there is no MAC entry in DMAC table (by sending unidirectional traffic), traffic is flooded and the input traffic is accounted in IMQ. For bidirectional traffic (MAC entry in DMAC table), if the outgoing interface is on the same I-chip then both ingress and egress statistics are counted in a combined way. If the outgoing interface is on a different I-chip or FPC, then only egress statistics are accounted in IMQ. This behavior is expected with non-enhanced DPCs</p>	detail extensive
Queue counters (Egress)	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	detail extensive
Ingress queues	<p>Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.</p>	extensive
Queue counters (Ingress)	<p>CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	extensive
Active alarms and Active defects	<p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the 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.</p> <ul style="list-style-type: none"> • None—There are no active defects or alarms. • Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. 	detail extensive none
OTN alarms	<p>Active OTN alarms identified on the interface.</p>	detail extensive

OTN defects	OTN defects received on the interface.	detail extensive
OTN FEC Mode	<p>The FECmode configured on the interface.</p> <ul style="list-style-type: none"> • efec—Enhanced forward error correction (EFEC) is configured to detect and correct bit errors. • gfec—G.709 Forward error correction (GFEC) mode is configured to detect and correct bit errors. • none—FEC mode is not configured. 	detail extensive
OTN Rate	<p>OTN mode.</p> <ul style="list-style-type: none"> • fixed-stuff-bytes—Fixed stuff bytes 11.0957 Gbps. • no-fixed-stuff-bytes—No fixed stuff bytes 11.0491 Gbps. • pass-through—Enable OTN passthrough mode. • no-pass-through—Do not enable OTN passthrough mode. 	detail extensive
OTN Line Loopback	Status of the line loopback, if configured for the DWDM OTN PIC. Its value can be: enabled or disabled .	detail extensive
OTN FEC statistics	<p>The forward error correction (FEC) counters for the DWDM OTN PIC.</p> <ul style="list-style-type: none"> • Corrected Errors—The count of corrected errors in the last second. • Corrected Error Ratio—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits. 	detail extensive
OTN FEC alarms	<p>OTN FEC excessive or degraded error alarms triggered on the interface.</p> <ul style="list-style-type: none"> • FEC Degrade—OTU FEC Degrade defect. • FEC Excessive—OTU FEC Excessive Error defect. 	detail extensive
OTN OC	<p>OTN OC defects triggered on the interface.</p> <ul style="list-style-type: none"> • LOS—OC Loss of Signal defect. • LOF—OC Loss of Frame defect. • LOM—OC Loss of Multiframe defect. • Wavelength Lock—OC Wavelength Lock defect. 	detail extensive

OTN OTU	<p>OTN OTU defects detected on the interface</p> <ul style="list-style-type: none"> • AIS—OTN AIS alarm. • BDI—OTN OTU BDI alarm. • IAE—OTN OTU IAE alarm. • TTIM—OTN OTU TTIM alarm. • SF—OTN ODU bit error rate fault alarm. • SD—OTN ODU bit error rate defect alarm. • TCA-ES—OTN ODU ES threshold alarm. • TCA-SES—OTN ODU SES threshold alarm. • TCA-UAS—OTN ODU UAS threshold alarm. • TCA-BBE—OTN ODU BBE threshold alarm. • BIP—OTN ODU BIP threshold alarm. • BBE—OTN OTU BBE threshold alarm. • ES—OTN OTU ES threshold alarm. • SES—OTN OTU SES threshold alarm. • UAS—OTN OTU UAS threshold alarm. 	detail extensive
Received DAPI	Destination Access Port Interface (DAPI) from which the packets were received.	detail extensive
Received SAPI	Source Access Port Interface (SAPI) from which the packets were received.	detail extensive
Transmitted DAPI	Destination Access Port Interface (DAPI) to which the packets were transmitted.	detail extensive
Transmitted SAPI	Source Access Port Interface (SAPI) to which the packets were transmitted.	detail extensive
PCS statistics	<p>(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device.</p> <ul style="list-style-type: none"> • Bit errors—The number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode. • Errored blocks—The number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode. 	detail extensive

MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> • Total octets and total packets—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. • 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 runs (which are normal occurrences caused by collisions) and noise hits are counted. • VLAN tagged frames—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. • Code violations—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error." 	extensive
OTN Received Overhead Bytes	APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08	extensive
OTN Transmitted Overhead Bytes	APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08	extensive

Filter statistics	<p>Receive and Transmit statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected.</p> <ul style="list-style-type: none"> • Input packet count—Number of packets received from the MAC hardware that the filter processed. • Input packet rejects—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • Input DA rejects—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the 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	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. 	extensive

WIS section	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B1—Bit interleaved parity for SONET section overhead • SEF—Severely errored framing • LOL—Loss of light • LOF—Loss of frame • ES-S—Errored seconds (section) • SES-S—Severely errored seconds (section) • SEFS-S—Severely errored framing seconds (section) 	extensive
WIS line	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. State other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B2—Bit interleaved parity for SONET line overhead • REI-L—Remote error indication (near-end line) • RDI-L—Remote defect indication (near-end line) • AIS-L—Alarm indication signal (near-end line) • BERR-SF—Bit error rate fault (signal failure) • BERR-SD—Bit error rate defect (signal degradation) • ES-L—Errored seconds (near-end line) • SES-L—Severely errored seconds (near-end line) • UAS-L—Unavailable seconds (near-end line) • ES-LFE—Errored seconds (far-end line) • SES-LFE—Severely errored seconds (far-end line) • UAS-LFE—Unavailable seconds (far-end line) 	extensive

WIS path	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none">• Seconds—Number of seconds the defect has been active.• Count—Number of times that the defect has gone from inactive to active.• State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none">• BIP-B3—Bit interleaved parity for SONET section overhead• REI-P—Remote error indication• LOP-P—Loss of pointer (path)• AIS-P—Path alarm indication signal• RDI-P—Path remote defect indication• UNEQ-P—Path unequipped• PLM-P—Path payload label mismatch• ES-P—Errored seconds (near-end STS path)• SES-P—Severely errored seconds (near-end STS path)• UAS-P—Unavailable seconds (near-end STS path)• SES-PFE—Severely errored seconds (far-end STS path)• UAS-PFE—Unavailable seconds (far-end STS path)	extensive
----------	--	-----------

Autonegotiation information	<p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> • Negotiation status: <ul style="list-style-type: none"> • Incomplete—Ethernet interface has the speed or link mode configured. • No autonegotiation—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation. • Complete—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner status—OK when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner: <ul style="list-style-type: none"> • Link mode—Depending on the capability of the attached Ethernet device, either Full-duplex or Half-duplex. • Flow control—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is None. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports both PAUSE on receive and transmit or only PAUSE receive). • Remote fault—Remote fault information from the link partner—Failure indicates a receive link error. OK indicates that the link partner is receiving. Negotiation error indicates a negotiation error. Offline indicates that the link partner is going offline. • Local resolution—Information from the link partner: <ul style="list-style-type: none"> • Flow control—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports both PAUSE on receive and transmit or only PAUSE receive). • Remote fault—Remote fault information. Link OK (no error detected on receive), Offline (local interface is offline), and Link Failure (link error detected on receive). 	extensive
Received path trace, Transmitted path trace	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the routing device at the other end of the fiber. The transmitted path trace value is the message that this routing device transmits.</p>	extensive

Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: <ul style="list-style-type: none"> • Destination slot—FPC slot number. 	extensive
CoS information	Information about the CoS queue for the physical interface. <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive

Logical Interface

Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> .	All levels

VLAN-Tag	<p>Rewrite profile applied to incoming or outgoing frames on the outer (Out) VLAN tag or for both the outer and inner (In) VLAN tags.</p> <ul style="list-style-type: none"> • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • pop—The outer VLAN tag of the incoming frame is removed. • swap—The outer VLAN tag of the incoming frame is overwritten with the user specified VLAN tag information. • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • push-push—Two VLAN tags are pushed in from the incoming frame. • swap-push—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • swap-swap—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user specified VLAN tag value. • pop-swap—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • pop-pop—Both the outer and inner VLAN tags of the incoming frame are removed. 	brief detail extensive none
Demux:	<p>IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following:</p> <ul style="list-style-type: none"> • Source Family Inet • Destination Family Inet 	detail extensive none
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under <i>Common Output Fields Description</i> .	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	detail extensive none

Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Input packets, Output packets—Number of packets received and transmitted on the interface set. 	detail extensive
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	extensive
Local statistics	Number and rate of bytes and packets destined to the routing device.	extensive
Transit statistics	<p>Number and rate of bytes and packets transiting the switch.</p> <p>NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the Output bytes and Output packets interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.</p>	extensive
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive none
Flags	Information about protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	detail extensive
Donor interface	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	detail extensive none
Preferred source address	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	detail extensive none
Input Filters	Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	detail extensive

Output Filters	Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	detail extensive
Mac-Validate Failures	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	detail extensive none
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	detail extensive none
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
Flags	Information about address flag (possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	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

For Gigabit Ethernet IQ PICs, traffic and MAC statistics output varies. The following table describes the traffic and MAC statistics for two sample interfaces, each of which is sending traffic in packets of 500 bytes (including 478 bytes for the Layer 3 packet, 18 bytes for the Layer 2 VLAN traffic header, and 4 bytes for cyclic redundancy check [CRC] information). The **ge-0/3/0** interface is the inbound physical interface, and the **ge-0/0/0** interface is the outbound physical interface. On both interfaces, traffic is carried on logical unit **.50** (VLAN 50).

Table 28: Gigabit and 10 Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type

Interface Type	Sample Command	Byte and Octet Counts Include	Comments
Inbound physical interface	show interfaces ge-0/3/0 extensive	Traffic statistics: Input bytes: 496 bytes per packet, representing the Layer 2 packet MAC statistics: Received octets: 500 bytes per packet, representing the Layer 2 packet + 4 bytes	The additional 4 bytes are for the CRC.
Inbound logical interface	show interfaces ge-0/3/0.50 extensive	Traffic statistics: Input bytes: 478 bytes per packet, representing the Layer 3 packet	
Outbound physical interface	show interfaces ge-0/0/0 extensive	Traffic statistics: Input bytes: 490 bytes per packet, representing the Layer 3 packet + 12 bytes MAC statistics: Received octets: 478 bytes per packet, representing the Layer 3 packet	For input bytes, the additional 12 bytes include 6 bytes for the destination MAC address plus 4 bytes for VLAN plus 2 bytes for the Ethernet type.
Outbound logical interface	show interfaces ge-0/0/0.50 extensive	Traffic statistics: Input bytes: 478 bytes per packet, representing the Layer 3 packet	

Table 29 on page 691 lists the output fields for the **show interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 29: show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Link-level type	Encapsulation being used on the physical interface.	All levels
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
MTU	Maximum transmission unit size on the physical interface.	All levels
Link mode	Link mode: Full-duplex or Half-duplex.	
Speed	Speed at which the interface is running.	All levels
BPDU error	Bridge protocol data unit (BPDU) error: Detected or None	
Loopback	Loopback status: Enabled or Disabled . If loopback is enabled, type of loopback: Local or Remote .	All levels
Source filtering	Source filtering status: Enabled or Disabled .	All levels
Flow control	Flow control status: Enabled or Disabled .	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: Enabled or Disabled .	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> • Online—Autonegotiation is manually configured as online. • Offline—Autonegotiation is manually configured as offline. 	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Link flags	Information about the physical link.	All levels

Table 29: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
CoS queues	Number of CoS queues configured.	detail extensive none
Current address	Configured MAC address.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None
Output Rate	Output rate in bps and pps.	None
Active alarms and Active defects	<p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. These fields can contain the value None or Link.</p> <ul style="list-style-type: none"> • None—There are no active defects or alarms. • Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. 	detail extensive none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive

Table 29: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Input errors	<p>Input errors on the interface.</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the ignore-l3-incompletes. • 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 29: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output errors	<p>Output errors on the interface.</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation; therefore, for Gigabit Ethernet PICs, this number must always remain 0. If it is nonzero, there is a software bug. • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field must never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive
Ingress queues	Total number of ingress queues supported on the specified interface.	extensive
Queue counters and queue number	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	detail extensive

Table 29: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> • Total octets and total packets—Total number of octets and packets. • Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets. • CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • FIFO error—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • MAC control frames—Number of MAC control frames. • MAC pause frames—Number of MAC control frames with pause operational code. • Oversized frames—There are two possible conditions regarding the number of oversized frames: <ul style="list-style-type: none"> • Packet length exceeds 1518 octets, or • Packet length exceeds MRU • Jabber frames—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • Fragment frames—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • VLAN tagged frames—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. • Code violations—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error." 	extensive

Table 29: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Filter statistics	<p>Receive and Transmit statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected.</p> <ul style="list-style-type: none"> • Input packet count—Number of packets received from the MAC hardware that the filter processed. • Input packet rejects—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • Input DA rejects—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the device from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local device (which the router is rejecting). • Input SA rejects—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • Output packet count—Number of packets that the filter has given to the MAC hardware. • Output packet pad count—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • Output packet error count—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment. • CAM destination filters, CAM source filters—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields must be 0. 	extensive

Table 29: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Autonegotiation information	<p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> • Negotiation status: <ul style="list-style-type: none"> • Incomplete—Ethernet interface has the speed or link mode configured. • No autonegotiation—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation. • Complete—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. 	extensive
Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number. 	extensive
CoS information	<p>Information about the CoS queue for the physical interface.</p> <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive
Interface transmit statistics	<p>Status of the interface-transmit-statistics configuration: Enabled or Disabled.</p>	detail extensive
Queue counters (Egress)	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	detail extensive

Table 29: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Input packets, Output packets—Number of packets received and transmitted on the interface set. 	detail extensive
Local statistics	Number and rate of bytes and packets destined to the device.	extensive
Transit statistics	<p>Number and rate of bytes and packets transiting the switch.</p> <p>NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the Output bytes and Output packets interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.</p>	extensive
Security	Security zones that interface belongs to.	extensive
Flow Input statistics	Statistics on packets received by flow module.	extensive

Table 29: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Flow Output statistics	Statistics on packets sent by flow module.	extensive
Flow error statistics (Packets dropped due to)	Statistics on errors in the flow module.	extensive
Protocol	Protocol family.	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive none
Flags	Information about protocol family flags. .	detail extensive
Addresses, Flags	Information about the address flags..	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Sample Output Gigabit Ethernet

show interfaces (Gigabit Ethernet)

```
user@host> show interfaces ge-3/0/2
```

```
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Interface index: 167, SNMP ifIndex: 35
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled
  Remote fault: Online
```

```

Device flags      : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
CoS queues       : 4 supported, 4 maximum usable queues
Current address: 00:00:5e:00:53:7c, Hardware address: 00:00:5e:00:53:7c
Last flapped    : 2006-08-10 17:25:10 PDT (00:01:08 ago)
Input rate       : 0 bps (0 pps)
Output rate      : 0 bps (0 pps)
Ingress rate at Packet Forwarding Engine      : 0 bps (0 pps)
Ingress drop rate at Packet Forwarding Engine : 0 bps (0 pps)
Active alarms    : None
Active defects   : None

Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69)
  Flags: SNMP-Traps 0x4000
  VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push
    0x8100.512 0x8100.513)
  Encapsulation: VLAN-CCC
  Egress account overhead: 100
  Ingress account overhead: 90
  Input packets : 0
  Output packets: 0
  Protocol ccc, MTU: 1522
  Flags: Is-Primary

```

show interfaces (Gigabit Ethernet on MX Series Routers)

user@host> **show interfaces ge-2/2/2**

```

Physical interface: ge-2/2/2, Enabled, Physical link is Up
  Interface index: 156, SNMP ifIndex: 188
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, MAC-REWRITE Error: None,
  Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags        : None
  CoS queues       : 8 supported, 4 maximum usable queues
  Schedulers        : 0
  Current address: 00:00:5e:00:53:c0, Hardware address: 00:00:5e:00:53:76
  Last flapped    : 2008-09-05 16:44:30 PDT (3d 01:04 ago)
  Input rate       : 0 bps (0 pps)
  Output rate      : 0 bps (0 pps)
  Active alarms    : None

```



```

Active defects : None
Logical interface ge-2/2/2.0 (Index 82) (SNMP ifIndex 219)
  Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
  Input packets : 10232
  Output packets: 10294
  Protocol inet, MTU: 1500
    Flags: Sendbroadcast-pkt-to-re
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 203.0.113/24, Local: 203.0.113.1, Broadcast: 203.0.113.255
  Protocol inet6, MTU: 1500
    Max nh cache: 4, New hold nh limit: 100000, Curr nh cnt: 4, Curr new hold cnt:
4, NH drop cnt: 0
    Flags: Is-Primary
    Addresses, Flags: Is-Default Is-Preferred Is-Primary
      Destination: 2001:db8:/32, Local: 2001:db8::5
    Addresses, Flags: Is-Preferred
      Destination: 2001:db8:1::/32, Local: 2001:db8:223:9cff:fe9f:3e78
  Protocol multiservice, MTU: Unlimited
    Flags: Is-Primary

```

show interfaces (link degrade status)

user@host> show interfaces et-3/0/0

```

Physical interface: et-3/0/0, Enabled, Physical link is Down
  Interface index: 157, SNMP ifIndex: 537
  Link-level type: Ethernet, MTU: 1514, MRU: 0, Speed: 100Gbps, BPDU Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 54:e0:32:23:9d:38, Hardware address: 54:e0:32:23:9d:38
  Last flapped   : 2014-06-18 02:36:38 PDT (02:50:50 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK
  PCS statistics
    Bit errors           Seconds
    Errored blocks       0
  Link Degraded* :
  Link Monitoring      : Enable

```

```

Link Degrade Set Threshold:      : 1E-7
Link Degrade Clear Threshold:    : 1E-12
Estimated BER                   : 1E-7
Link-degrade event              : Seconds      Count      State
                                782             1        Defect Active

```

show interfaces extensive (Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration)

```
user@host> show interfaces ge-2/1/2 extensive | match "output|interface"
```

```

Physical interface: ge-2/1/2, Enabled, Physical link is Up
Interface index: 151, SNMP ifIndex: 530, Generation: 154
Interface flags: SNMP-Traps Internal: 0x4000
Output bytes      :          240614363944          772721536 bps
Output packets:    3538446506             1420444 pps
Direction : Output
Interface transmit statistics: Enabled

Logical interface ge-2/1/2.0 (Index 331) (SNMP ifIndex 955) (Generation 146)
Output bytes      :          195560312716          522726272 bps
Output packets:    4251311146             1420451 pps

```

```
user@host> show interfaces ge-5/2/0.0 statistics detail
```

```

Logical interface ge-5/2/0.0 (Index 71) (SNMP ifIndex 573) (Generation 135)
Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
Egress account overhead: 100
Ingress account overhead: 90
Traffic statistics:
Input bytes      :          271524
Output bytes     :        37769598
Input packets:    3664
Output packets:   885790
IPv6 transit statistics:
Input bytes      :          0
Output bytes     :        16681118
Input packets:    0
Output packets:   362633
Local statistics:
Input bytes      :          271524
Output bytes     :        308560
Input packets:    3664

```

```

Output packets:                3659
Transit statistics:
Input  bytes   :                0          0 bps
Output bytes   :            37461038        0 bps
Input  packets :                0          0 pps
Output packets :            882131         0 pps
IPv6 transit statistics:
Input  bytes   :                0          0 bps
Output bytes   :            16681118        0 bps
Input  packets :                0          0 pps
Output packets :            362633         0 pps

```

show interfaces brief (Gigabit Ethernet)

user@host> **show interfaces ge-3/0/2 brief**

```

Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags      : Present Running
  Interface flags:  SNMP-Traps Internal: 0x4000
  Link flags        : None

Logical interface ge-3/0/2.0
  Flags: SNMP-Traps 0x4000
  VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push
  0x8100.512 0x8100.513)
  Encapsulation: VLAN-CCC
  ccc

Logical interface ge-3/0/2.32767
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ] Encapsulation: ENET2

```

show interfaces detail (Gigabit Ethernet)

user@host> **show interfaces ge-3/0/2 detail**

```

Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Interface index: 167, SNMP ifIndex: 35, Generation: 177
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online

```

```

Device flags      : Present Running
Interface flags:  SNMP-Traps Internal: 0x4000
Link flags       : None
CoS queues       : 4 supported, 4 maximum usable queues
Hold-times       : Up 0 ms, Down 0 ms
Current address:  00:00:5e:00:53:7c, Hardware address: 00:00:5e:00:53:7c
Last flapped    : 2006-08-09 17:17:00 PDT (01:31:33 ago)
Statistics last cleared: Never

Traffic statistics:
  Input  bytes :                0                0 bps
  Output bytes :                0                0 bps
  Input  packets:                0                0 pps
  Output packets:                0                0 pps
Ingress traffic statistics at Packet Forwarding Engine:
  Input  bytes :                0                0 bps
  Input  packets:                0                0 pps
  Drop  bytes  :                0                0 bps
  Drop  packets:                0                0 pps
Ingress queues: 4 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets
  0 best-effort      0                0                0
  1 expedited-fo     0                0                0
  2 assured-forw     0                0                0
  3 network-cont     0                0                0
Egress queues: 4 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets
  0 best-effort      0                0                0
  1 expedited-fo     0                0                0
  2 assured-forw     0                0                0
  3 network-cont     0                0                0
Active alarms  : None
Active defects : None

Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69) (Generation 140)
  Flags: SNMP-Traps 0x4000
  VLAN-Tag [0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530)
Out(swap-push 0x8100.512 0x8100.513)
  Encapsulation: VLAN-CCC
  Egress account overhead: 100
  Ingress account overhead: 90
  Traffic statistics:
    Input  bytes :                0
    Output bytes :                0
    Input  packets:                0

```

```

    Output packets:                0
Local statistics:
    Input  bytes   :                0
    Output bytes   :                0
    Input  packets:                0
    Output packets:                0
Transit statistics:
    Input  bytes   :                0          0 bps
    Output bytes   :                0          0 bps
    Input  packets:                0          0 pps
    Output packets:                0          0 pps
Protocol ccc, MTU: 1522, Generation: 149, Route table: 0
  Flags: Is-Primary

Logical interface ge-3/0/2.32767 (Index 71) (SNMP ifIndex 70)
(Generation 139)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
Traffic statistics:
    Input  bytes   :                0
    Output bytes   :                0
    Input  packets:                0
    Output packets:                0
Local statistics:
    Input  bytes   :                0
    Output bytes   :                0
    Input  packets:                0
    Output packets:                0
Transit statistics:
    Input  bytes   :                0          0 bps
    Output bytes   :                0          0 bps
    Input  packets:                0          0 pps
    Output packets:                0          0 pps

```

show interfaces extensive (Gigabit Ethernet IQ2)

user@host> **show interfaces ge-7/1/3 extensive**

```

Physical interface: ge-7/1/3, Enabled, Physical link is Up
Interface index: 170, SNMP ifIndex: 70, Generation: 171
Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x4004000

```

```

Link flags      : None
CoS queues      : 8 supported, 4 maximum usable queues
Schedulers     : 256
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:00:5e:00:53:74, Hardware address: 00:00:5e:00:53:74
Last flapped   : 2007-11-07 21:31:41 PST (02:03:33 ago)
Statistics last cleared: Never
Traffic statistics:
  Input bytes   :          38910844056          7952 bps
  Output bytes  :          7174605          8464 bps
  Input packets:          418398473          11 pps
  Output packets:          78903          12 pps
IPv6 transit statistics:
  Input bytes   :          0
  Output bytes  :          0
  Input packets:          0
  Output packets:          0
Ingress traffic statistics at Packet Forwarding Engine:
  Input bytes   :          38910799145          7952 bps
  Input packets:          418397956          11 pps
  Drop bytes    :          0          0 bps
  Drop packets  :          0          0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Ingress queues: 4 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets  Dropped packets
  0 best-effort      418390823      418390823          0
  1 expedited-fo          0          0          0
  2 assured-forw          0          0          0
  3 network-cont      7133      7133          0
Egress queues: 4 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets  Dropped packets
  0 best-effort      1031      1031          0
  1 expedited-fo          0          0          0
  2 assured-forw          0          0          0
  3 network-cont      77872      77872          0
Active alarms : None
Active defects : None
MAC statistics:      Receive      Transmit

```

```

Total octets                38910844056                7174605
Total packets                418398473                78903
Unicast packets              408021893366                1026
Broadcast packets            10                        12
Multicast packets            418398217                77865
CRC/Align errors             0                        0
FIFO errors                  0                        0
MAC control frames           0                        0
MAC pause frames             0                        0
Oversized frames             0
Jabber frames                0
Fragment frames              0
VLAN tagged frames           0
Code violations               0   OTN Received Overhead Bytes:
APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58
Payload Type: 0x08
OTN Transmitted Overhead Bytes:
APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
Payload Type: 0x08
Filter statistics:
Input packet count           418398473
Input packet rejects          479
Input DA rejects              479
Input SA rejects              0
Output packet count           78903
Output packet pad count       0
Output packet error count     0
CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
Negotiation status: Complete
Link partner:
Link mode: Full-duplex, Flow control: Symmetric/Asymmetric,
Remote fault: OK
Local resolution:
Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
Destination slot: 7
CoS information:
Direction : Output
CoS transmit queue           Bandwidth           Buffer           Priority           Limit
                                %                bps                %                usec
0 best-effort                 95                950000000        95                0                low        none
3 network-control             5                  50000000         5                 0                low        none

```

Direction : Input

CoS transmit queue	Bandwidth		Buffer		Priority	Limit
	%	bps	%	usec		
0 best-effort	95	950000000	95	0	low	none
3 network-control	5	50000000	5	0	low	none

Logical interface ge-7/1/3.0 (Index 70) (SNMP ifIndex 85) (Generation 150)

Flags: SNMP-Traps Encapsulation: ENET2

Traffic statistics:

```

Input  bytes :      812400
Output bytes :     1349206
Input  packets:      9429
Output packets:     9449

```

IPv6 transit statistics:

```

Input  bytes :      0
Output bytes :      0
Input  packets:      0
Output packets:      0

```

Local statistics:

```

Input  bytes :      812400
Output bytes :     1349206
Input  packets:      9429
Output packets:     9449

```

Transit statistics:

```

Input  bytes :      0          7440 bps
Output bytes :      0          7888 bps
Input  packets:      0          10 pps
Output packets:      0          11 pps

```

IPv6 transit statistics:

```

Input  bytes :      0
Output bytes :      0
Input  packets:      0
Output packets:      0

```

Protocol inet, MTU: 1500, Generation: 169, Route table: 0

Flags: Is-Primary, Mac-Validate-Strict

Mac-Validate Failures: Packets: 0, Bytes: 0

Addresses, Flags: Is-Preferred Is-Primary

Input Filters: F1-ge-3/0/1.0-in, F3-ge-3/0/1.0-in

Output Filters: F2-ge-3/0/1.0-out (53)

Destination: 203.0.113/24, Local: 203.0.113.2, Broadcast: 203.0.113.255,
Generation: 196


```
Protocol multiservice, MTU: Unlimited, Generation: 170, Route table: 0
Flags: Is-Primary
Policer: Input: __default_arp_policer__
```

NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics displayed in the **show interfaces** command output might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the interface counters. For detailed information, see the description of the logical interface **Transit statistics** fields in [Table 27 on page 652](#).

show interfaces (Gigabit Ethernet Unnumbered Interface)

```
user@host> show interfaces ge-3/2/0
```

```
Physical interface: ge-3/2/0, Enabled, Physical link is Up
  Interface index: 148, SNMP ifIndex: 50
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Current address: 00:00:5e:00:53:f8, Hardware address: 00:00:5e:00:53:f8
  Last flapped   : 2006-10-27 04:42:23 PDT (08:01:52 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 624 bps (1 pps)
  Active alarms  : None
  Active defects : None

Logical interface ge-3/2/0.0 (Index 67) (SNMP ifIndex 85)
  Flags: SNMP-Traps Encapsulation: ENET2
  Input packets : 0
  Output packets: 6
  Protocol inet, MTU: 1500
    Flags: Unnumbered
    Donor interface: lo0.0 (Index 64)
    Preferred source address: 203.0.113.22
```

show interfaces (ACI Interface Set Configured)

```
user@host> show interfaces ge-1/0/0.4001
```

```

Logical interface ge-1/0/0.4001 (Index 340) (SNMP ifIndex 548)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.4001 ]  Encapsulation: PPP-over-
Ethernet
ACI VLAN:
  Dynamic Profile: aci-vlan-set-profile
PPPoE:
  Dynamic Profile: aci-vlan-pppoe-profile,
  Service Name Table: None,
  Max Sessions: 32000, Max Sessions VSA Ignore: Off,
  Duplicate Protection: On, Short Cycle Protection: Off,
  Direct Connect: Off,
  AC Name: nbc
Input packets : 9
Output packets: 8
Protocol multiservice, MTU: Unlimited

```

show interfaces (All Interface Set)

```
user@host> show interfaces ge-1/0/0.10
```

```

Logical interface ge-1/0/0.10 (Index 346) (SNMP ifIndex 554) (Generation 155)
  Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.10 ]  Encapsulation: ENET2
Line Identity:
  Dynamic Profile: ali-set-profile
  Circuit-id Remote-id Accept-no-ids
PPPoE:
  Dynamic Profile: ali-vlan-pppoe-profile,
  Service Name Table: None,
  Max Sessions: 32000, Max Sessions VSA Ignore: Off,
  Duplicate Protection: On, Short Cycle Protection: Off,
  Direct Connect: Off,
  AC Name: nbc
Input packets : 9
Output packets: 8
Protocol multiservice, MTU: Unlimited

```

Sample Output Gigabit Ethernet

```
show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, IQ2)
```

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

```

Physical interface: xe-5/0/0, Enabled, Physical link is Up
  Interface index: 177, SNMP ifIndex: 630, Generation: 178
  Link-level type: Ethernet, MTU: 1518, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Enabled,
  Flow control: Enabled
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags       : None
  CoS queues       : 8 supported, 4 maximum usable queues
  Schedulers       : 1024
  Hold-times       : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:f6, Hardware address: 00:00:5e:00:53:f6
  Last flapped    : Never
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   :          6970332384          0 bps
    Output bytes  :              0          0 bps
    Input packets :          81050506          0 pps
    Output packets:              0          0 pps
  IPv6 transit statistics:
    Input bytes   :              0
    Output bytes  :              0
    Input packets :              0
    Output packets:              0
  Ingress traffic statistics at Packet Forwarding Engine:
    Input bytes   :          6970299398          0 bps
    Input packets :          81049992          0 pps
    Drop bytes    :              0          0 bps
    Drop packets  :              0          0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0,
    L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0,
    MTU errors: 0, Resource errors: 0
  Ingress queues: 4 supported, 4 in use
  Queue counters:
    Queued packets  Transmitted packets  Dropped packets
    0 best-effort    81049992          81049992          0
    1 expedited-fo    0              0              0
    2 assured-forw    0              0              0
    3 network-cont    0              0              0
  Egress queues: 4 supported, 4 in use

```

```

Queue counters:      Queued packets  Transmitted packets      Dropped packets
  0 best-effort      0                0                0
  1 expedited-fo     0                0                0
  2 assured-forw     0                0                0
  3 network-cont     0                0                0
Active alarms  : None
Active defects : None
PCS statistics      Seconds
  Bit errors        0
  Errored blocks    0
MAC statistics:      Receive      Transmit
  Total octets      6970332384      0
  Total packets     81050506       0
  Unicast packets   81050000       0
  Broadcast packets 506          0
  Multicast packets 0            0
  CRC/Align errors  0            0
  FIFO errors       0            0
  MAC control frames 0            0
  MAC pause frames  0            0
  Oversized frames  0
  Jabber frames     0
  Fragment frames   0
  VLAN tagged frames 0
  Code violations    0
Filter statistics:
  Input packet count 81050506
  Input packet rejects 506
  Input DA rejects    0
  Input SA rejects    0
  Output packet count 0
  Output packet pad count 0
  Output packet error count 0
  CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 5
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority  Limit
                           %      bps      %      usec
  0 best-effort           95      950000000    95      0      low      none
  3 network-control       5       50000000     5      0      low      none
  Direction : Input
  CoS transmit queue      Bandwidth      Buffer Priority  Limit

```

```

          %          bps      %          usec
0 best-effort      95      950000000    95          0          low      none
3 network-control   5       50000000     5          0          low      none

Logical interface xe-5/0/0.0 (Index 71) (SNMP ifIndex 95) (Generation 195)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ] Encapsulation: ENET2
  Egress account overhead: 100
  Ingress account overhead: 90
  Traffic statistics:
    Input  bytes   :                0
    Output bytes   :               46
    Input  packets :                0
    Output packets :                1
  IPv6 transit statistics:
    Input  bytes   :                0
    Output bytes   :                0
    Input  packets :                0
    Output packets :                0
  Local statistics:
    Input  bytes   :                0
    Output bytes   :               46
    Input  packets :                0
    Output packets :                1
  Transit statistics:
    Input  bytes   :                0          0 bps
    Output bytes   :                0          0 bps
    Input  packets :                0          0 pps
    Output packets :                0          0 pps
  IPv6 transit statistics:
    Input  bytes   :                0
    Output bytes   :                0
    Input  packets :                0
    Output packets :                0
  Protocol inet, MTU: 1500, Generation: 253, Route table: 0
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 192.0.2/24, Local: 192.0.2.1, Broadcast: 192.0.2.255,
  Generation: 265
  Protocol multiservice, MTU: Unlimited, Generation: 254, Route table: 0
    Flags: None
    Policers: Input: __default_arp_policer__

```

show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode)

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

```

Physical interface: xe-1/0/0, Enabled, Physical link is Up
  Interface index: 141, SNMP ifIndex: 630, Generation: 47
  Link-level type: Ethernet, MTU: 1514, Speed: 9.294GbpsGbps, Loopback: Disabled
  WAN-PHY mode
  Source filtering: Disabled, Flow control: Enabled Speed Configuration: Auto
  Device flags      : Present Running
  Interface flags: SNMP-Traps 16384
  Link flags        : None
  CoS queues        : 4 supported
  Hold-times        : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:9d, Hardware address: 00:00:5e:00:53:9d
  Last flapped     : 2005-07-07 11:22:34 PDT (3d 12:28 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input  bytes :                0                0 bps
    Output bytes :                0                0 bps
    Input  packets:                0                0 pps
    Output packets:                0                0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    HS Link CRC errors: 0, HS Link FIFO overflows: 0,
    Resource errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0,
    Aged packets: 0, FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0,
    Resource errors: 0
  Queue counters:
    Queued packets  Transmitted packets  Dropped packets
    0 best-effort    0                0                0
    1 expedited-fo   0                0                0
    2 assured-forw    0                0                0
    3 network-cont    0                0                0
  Active alarms : LOL, LOS, LBL
  Active defects: LOL, LOS, LBL, SEF, AIS-L, AIS-P
  PCS statistics
    Seconds  Count
    Bit errors      0      0
    Errored blocks  0      0
  MAC statistics:
    Receive  Transmit
    Total octets      0      0
    Total packets     0      0
    Unicast packets   0      0
    Broadcast packets 0      0
    Multicast packets 0      0
    CRC/Align errors  0      0

```

```

FIFO errors                                0          0
MAC control frames                        0          0
MAC pause frames                         0          0
Oversized frames                        0
Jabber frames                           0
Fragment frames                         0
VLAN tagged frames                      0
Code violations                          0
Filter statistics:
  Input packet count                     0
  Input packet rejects                   0
  Input DA rejects                       0
  Input SA rejects                       0
  Output packet count                    0
  Output packet pad count                0
  Output packet error count              0
  CAM destination filters: 0, CAM source filters: 0
PMA PHY:          Seconds      Count  State
  PLL lock         0           0  OK
  PHY light        63159       1  Light Missing
WIS section:
  BIP-B1           0           0
  SEF              434430      434438  Defect Active
  LOS              434430       1  Defect Active
  LOF              434430       1  Defect Active
  ES-S             434430
  SES-S            434430
  SEFS-S           434430
WIS line:
  BIP-B2           0           0
  REI-L            0           0
  RDI-L            0           0  OK
  AIS-L            434430       1  Defect Active
  BERR-SF          0           0  OK
  BERR-SD          0           0  OK
  ES-L             434430
  SES-L            434430
  UAS-L            434420
  ES-LFE           0
  SES-LFE          0
  UAS-LFE          0
WIS path:
  BIP-B3           0           0
  REI-P            0           0

```

```

LOP-P          0          0 OK
AIS-P          434430      1 Defect Active
RDI-P          0          0 OK
UNEQ-P         0          0 OK
PLM-P          0          0 OK
ES-P           434430
SES-P           434430
UAS-P           434420
ES-PFE         0
SES-PFE         0
UAS-PFE         0

```

Received path trace:

```
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
```

Transmitted path trace: orissa so-1/0/0

```
6f 72 69 73 73 61 20 73 6f 2d 31 2f 30 2f 30 00 orissa so-1/0/0.
```

Packet Forwarding Engine configuration:

Destination slot: 1

CoS information:

CoS transmit queue		Bandwidth		Buffer	Priority	Limit
	%	bps	%	bytes		
0 best-effort	95	950000000	95	0	low	none
3 network-control	5	50000000	5	0	low	none

show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC)

user@host> show interfaces ge-7/0/0 extensive

```

Physical interface: ge-7/0/0, Enabled, Physical link is Down
Interface index: 143, SNMP ifIndex: 508, Generation: 208
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
Flow control: Enabled
Device flags   : Present Running Down
Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
Link flags     : None
Wavelength    : 1550.12 nm, Frequency: 193.40 THz
CoS queues     : 8 supported, 8 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:00:5e:00:53:72, Hardware address: 00:00:5e:00:53:72
Last flapped   : 2011-04-20 15:48:54 PDT (18:39:49 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes   : 0          0 bps
Output bytes  : 0          0 bps

```



```

Input  packets:                0                0 pps
Output packets:                0                0 pps
IPv6 transit statistics:
  Input  bytes  :                0
  Output bytes  :                0
  Input  packets:                0
  Output packets:                0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 2, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets
  0 best-effort             0                0                0
  1 expedited-fo           0                0                0
  2 assured-forw           0                0                0
  3 network-cont
Queue number:      Mapped forwarding classes
  0                best-effort
  1                expedited-forwarding
  2                assured-forwarding
  3                network-control
Active alarms  : LINK
Active defects : LINK
MAC statistics:
Total octets      Receive      Transmit
Total packets      0                0
Unicast packets    0                0
Broadcast packets  0                0
Multicast packets  0                0
CRC/Align errors   0                0
FIFO errors        0                0
MAC control frames  0                0
MAC pause frames   0                0
Oversized frames   0
Jabber frames      0
Fragment frames    0
VLAN tagged frames  0
Code violations     0
Total octets      0                0
Total packets      0                0

```

```

Unicast packets                0                0
Broadcast packets              0                0
Multicast packets              0                0
CRC/Align errors               0                0
FIFO errors                    0                0
MAC control frames             0                0
MAC pause frames               0                0
Oversized frames               0
Jabber frames                  0
Fragment frames                0
VLAN tagged frames             0
Code violations                 0

OTN alarms                     :   None
OTN defects                     :   None
OTN FEC Mode                    : GFEC
OTN Rate                       : Fixed Stuff Bytes 11.0957Gbps
OTN Line Loopback : Enabled
OTN FEC statistics :
    Corrected Errors                                0
    Corrected Error Ratio (          0 sec average)  0e-0
OTN FEC alarms:          Seconds      Count  State
    FEC Degrade           0           0  OK
    FEC Excessive         0           0  OK
OTN OC:                  Seconds      Count  State
    LOS                   2           1  OK
    LOF                   67164       2  Defect Active
    LOM                   67164       71  Defect Active
    Wavelength Lock       0           0  OK
OTN OTU:
    AIS                   0           0  OK
    BDI                   65919      4814  Defect Active
    IAE                   67158       1  Defect Active
    TTIM                  7           1  OK
    SF                    67164       2  Defect Active
    SD                    67164       3  Defect Active
    TCA-ES                0           0  OK
    TCA-SES               0           0  OK
    TCA-UAS               80          40  OK
    TCA-BBE               0           0  OK
    BIP                   0           0  OK
    BBE                   0           0  OK
    ES                    0           0  OK
    SES                   0           0  OK
    UAS                   587         0  OK

```

```
Received DAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Received SAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Transmitted DAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Transmitted SAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
OTN Received Overhead Bytes:
  APS/PCC0: 0x02, APS/PCC1: 0x42, APS/PCC2: 0xa2, APS/PCC3: 0x48
  Payload Type: 0x03
OTN Transmitted Overhead Bytes:
  APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
  Payload Type: 0x03
Filter statistics:
  Input packet count                0
  Input packet rejects              0
  Input DA rejects                  0
  Input SA rejects                  0
  Output packet count                0
  Output packet pad count            0
  Output packet error count          0
  CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 7
CoS information:
  Direction : Output
  CoS transmit queue                Bandwidth                Buffer Priority
Limit
                                %      bps      %      usec
    0 best-effort                95    9500000000    95         0      low
none
    3 network-control            5     500000000     5         0      low
none
    ...
```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode)

user@host> **show interfaces xe-7/0/0 extensive**

```
Physical interface: xe-7/0/0, Enabled, Physical link is Up
  Interface index: 173, SNMP ifIndex: 212, Generation: 174
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Unidirectional:
  Enabled,
```

```

Loopback: None, Source filtering: Disabled, Flow control: Enabled
Device flags   : Present Running
...

```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only)

user@host> show interfaces xe-7/0/0-tx extensive

```

Physical interface: xe-7/0/0-tx, Enabled, Physical link is Up
  Interface index: 176, SNMP ifIndex: 137, Generation: 177
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Unidirectional:
  Tx-Only
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:83, Hardware address: 00:00:5e:00:53:83
  Last flapped   : 2007-06-01 09:08:19 PDT (3d 02:31 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input  bytes   :                0                0 bps
    Output bytes   :    322891152287160    9627472888 bps
    Input  packets :                0                0 pps
    Output packets :    328809727380    1225492 pps

...

  Filter statistics:
    Output packet count      328810554250
    Output packet pad count   0
    Output packet error count 0
...

Logical interface xe-7/0/0-tx.0 (Index 73) (SNMP ifIndex 138) (Generation 139)
  Flags: SNMP-Traps Encapsulation: ENET2
  Egress account overhead: 100
  Ingress account overhead: 90
  Traffic statistics:
    Input  bytes   :                0
    Output bytes   :    322891152287160
    Input  packets :                0
    Output packets :    328809727380
  IPv6 transit statistics:

```

```

    Input  bytes :                0
    Output bytes :                0
    Input  packets:                0
    Output packets:                0
Local statistics:
    Input  bytes :                0
    Output bytes :                0
    Input  packets:                0
    Output packets:                0
Transit statistics:
    Input  bytes :                0                0 bps
    Output bytes :      322891152287160      9627472888 bps
    Input  packets:                0                0 pps
    Output packets:      328809727380      1225492 pps
IPv6 transit statistics:
    Input  bytes :                0
    Output bytes :                0
    Input  packets:                0
    Output packets:                0
Protocol inet, MTU: 1500, Generation: 147, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.11.12/24, Local: 10.11.12.13, Broadcast: 10.11.12.255,
Generation: 141
  Protocol multiservice, MTU: Unlimited, Generation: 148, Route table: 0
  Flags: None
  Policer: Input: __default_arp_policer__

```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only)

user@host> show interfaces xe-7/0/0-rx extensive

```

Physical interface: xe-7/0/0-rx, Enabled, Physical link is Up
  Interface index: 174, SNMP ifIndex: 118, Generation: 175
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Unidirectional:
Rx-Only
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags        : None
  CoS queues        : 8 supported, 8 maximum usable queues
  Hold-times        : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:83, Hardware address: 00:00:5e:00:53:83
  Last flapped      : 2007-06-01 09:08:22 PDT (3d 02:31 ago)
  Statistics last cleared: Never
  Traffic statistics:

```

```

Input bytes :      322857456303482      9627496104 bps
Output bytes :      0      0 bps
Input packets:      328775413751      1225495 pps
Output packets:      0      0 pps

```

...

Filter statistics:

```

Input packet count      328775015056
Input packet rejects      1
Input DA rejects      0

```

...

Logical interface xe-7/0/0-rx.0 (Index 72) (SNMP ifIndex 120) (Generation 138)

Flags: SNMP-Traps Encapsulation: ENET2

Traffic statistics:

```

Input bytes :      322857456303482
Output bytes :      0
Input packets:      328775413751
Output packets:      0

```

IPv6 transit statistics:

```

Input bytes :      0
Output bytes :      0
Input packets:      0
Output packets:      0

```

Local statistics:

```

Input bytes :      0
Output bytes :      0
Input packets:      0
Output packets:      0

```

Transit statistics:

```

Input bytes :      322857456303482      9627496104 bps
Output bytes :      0      0 bps
Input packets:      328775413751      1225495 pps
Output packets:      0      0 pps

```

IPv6 transit statistics:

```

Input bytes :      0
Output bytes :      0
Input packets:      0
Output packets:      0

```

Protocol inet, MTU: 1500, Generation: 145, Route table: 0

Addresses, Flags: Is-Preferred Is-Primary

Destination: 192.0.2/24, Local: 192.0.2.1, Broadcast: 192.0.2.255,

```

Generation: 139
  Protocol multiservice, MTU: Unlimited, Generation: 146, Route table: 0
  Flags: None
  Policers: Input: __default_arp_policer__

```

Sample Output

Sample Output SRX Gigabit Ethernet

```
user@host> show interfaces ge-0/0/1
```

```

Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:42 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK
  Interface transmit statistics: Disabled

Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514)
  Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
  Security: Zone: public
  Protocol inet, MTU: 1500
    Flags: Sendbroadcast-pkt-to-re
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255

```

Sample Output SRX Gigabit Ethernet

```
user@host> show interfaces ge-0/0/1
```

```
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags      : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags        : None
  CoS queues        : 8 supported, 8 maximum usable queues
  Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
  Last flapped      : 2015-05-12 08:36:59 UTC (1w1d 22:42 ago)
  Input rate        : 0 bps (0 pps)
  Output rate       : 0 bps (0 pps)
  Active alarms     : LINK
  Active defects    : LINK
  Interface transmit statistics: Disabled

Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514)
  Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
  Security: Zone: public
  Protocol inet, MTU: 1500
    Flags: Sendbroadcast-pkt-to-re
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255
```

show interfaces (Gigabit Ethernet for vSRX and vSRX 3.0)

```
user@host> show interfaces ge-0/0/0
```

```
Physical interface: ge-0/0/0, Enabled, Physical link is Up
  Interface index: 136, SNMP ifIndex: 510
  Link-level type: Ethernet, MTU: 1518, LAN-PHY mode, Link-mode: Half-duplex,
  Speed: 1000mbps, BPDU Error: None, Loop Detect PDU Error: None, Ethernet-Switching
  Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
  Disabled, Flow control: Enabled,
  Auto-negotiation: Enabled, Remote fault: Online
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues        : 8 supported, 8 maximum usable queues
  Current address: 00:50:56:93:ef:25, Hardware address: 00:50:56:93:ef:25
  Last flapped      : 2019-03-29 01:57:45 UTC (00:00:41 ago)
```



```

Input rate      : 1120 bps (0 pps)
Output rate     : 0 bps (0 pps)
Active alarms   : None

```

show interfaces detail (Gigabit Ethernet)

user@host> show interfaces ge-0/0/1 detail

```

Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510, Generation: 138
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
  Disabled,
  Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w2d 00:00 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input  bytes :                0                0 bps
    Output bytes :                0                0 bps
    Input  packets:                0                0 pps
    Output packets:                0                0 pps
  Egress queues: 8 supported, 4 in use
  Queue counters:
    Queued packets  Transmitted packets  Dropped packets
    0 best-effort   0                    0                    0
    1 expedited-fo  0                    0                    0
    2 assured-forw  0                    0                    0
    3 network-cont  0                    0                    0
  Queue number:
    Mapped forwarding classes
    0 best-effort
    1 expedited-forwarding
    2 assured-forwarding
    3 network-control
  Active alarms   : LINK
  Active defects  : LINK
  Interface transmit statistics: Disabled

Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)
  Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2

```

```

Traffic statistics:
  Input  bytes   :          0
  Output bytes   :          0
  Input  packets:          0
  Output packets:          0
Local statistics:
  Input  bytes   :          0
  Output bytes   :          0
  Input  packets:          0
  Output packets:          0
Transit statistics:
  Input  bytes   :          0          0 bps
  Output bytes   :          0          0 bps
  Input  packets:          0          0 pps
  Output packets:          0          0 pps
Security: Zone: public
Flow Statistics :
Flow Input statistics :
  Self packets :          0
  ICMP packets :          0
  VPN packets  :          0
  Multicast packets :          0
  Bytes permitted by policy :          0
  Connections established :          0
Flow Output statistics:
  Multicast packets :          0
  Bytes permitted by policy :          0
Flow error statistics (Packets dropped due to):
  Address spoofing:          0
  Authentication failed:          0
  Incoming NAT errors:          0
  Invalid zone received packet:          0
  Multiple user authentications:          0
  Multiple incoming NAT:          0
  No parent for a gate:          0
  No one interested in self packets:          0
  No minor session:          0
  No more sessions:          0
  No NAT gate:          0
  No route present:          0
  No SA for incoming SPI:          0
  No tunnel found:          0
  No session for a gate:          0
  No zone or NULL zone binding          0

```

```

Policy denied: 0
Security association not active: 0
TCP sequence number out of window: 0
Syn-attack protection: 0
User authentication errors: 0
Protocol inet, MTU: 1500, Generation: 150, Route table: 0
Flags: Sendbroadcast-pkt-to-re
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255, Generation:
150

```

show interfaces statistics st0.0 detail

user@host> show interfaces statistics st0.0 detail

```

Logical interface st0.0 (Index 71) (SNMP ifIndex 609) (Generation 136)
Flags: Up Point-To-Point SNMP-Traps Encapsulation: Secure-Tunnel
Traffic statistics:
Input bytes : 528152756774
Output bytes : 575950643520
Input packets: 11481581669
Output packets: 12520666095
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 0 121859888 bps
Output bytes : 0 128104112 bps
Input packets: 0 331141 pps
Output packets: 0 348108 pps
Security: Zone: untrust
Allowed host-inbound traffic : any-service bfd bgp dvmrp igmp ldp msdp nhrp
ospf ospf3 pgm pim rip ripng router-discovery rsvp
sap vrrp
Flow Statistics :
Flow Input statistics :
Self packets : 0
ICMP packets : 0
VPN packets : 0
Multicast packets : 0
Bytes permitted by policy : 525984295844
Connections established : 7

```

```

Flow Output statistics:
  Multicast packets :          0
  Bytes permitted by policy :    576003290222
Flow error statistics (Packets dropped due to):
  Address spoofing:            0
  Authentication failed:       0
  Incoming NAT errors:         0
  Invalid zone received packet: 0
  Multiple user authentications: 0
  Multiple incoming NAT:       0
  No parent for a gate:        0
  No one interested in self packets: 0
  No minor session:            0
  No more sessions:            0
  No NAT gate:                 0
  No route present:            2000280
  No SA for incoming SPI:      0
  No tunnel found:             0
  No session for a gate:       0
  No zone or NULL zone binding 0
  Policy denied:               0
  Security association not active: 0
  TCP sequence number out of window: 0
  Syn-attack protection:       0
  User authentication errors:   0
Protocol inet, MTU: 9192
Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0,
NH drop cnt: 0
Generation: 155, Route table: 0
Flags: Sendbroadcast-pkt-to-re

```

show interfaces extensive (Gigabit Ethernet)

user@host> show interfaces ge-0/0/1.0 extensive

```

Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510, Generation: 138
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0

```

```

Link flags      : None
CoS queues     : 8 supported, 8 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:57 ago)
Statistics last cleared: Never
Traffic statistics:
  Input  bytes :                0                0 bps
  Output bytes :                0                0 bps
  Input packets:                0                0 pps
  Output packets:              0                0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets
  0 best-effort      0                0                0
  1 expedited-fo     0                0                0
  2 assured-forw     0                0                0
  3 network-cont     0                0                0
Queue number:      Mapped forwarding classes
  0                best-effort
  1                expedited-forwarding
  2                assured-forwarding
  3                network-control
Active alarms  : LINK
Active defects : LINK
MAC statistics:      Receive      Transmit
  Total octets      0                0
  Total packets      0                0
  Unicast packets    0                0
  Broadcast packets  0                0
  Multicast packets  0                0
  CRC/Align errors   0                0
  FIFO errors        0                0
  MAC control frames  0                0
  MAC pause frames   0                0
  Oversized frames   0
  Jabber frames      0
  Fragment frames    0

```

```

VLAN tagged frames          0
Code violations              0
Filter statistics:
  Input packet count        0
  Input packet rejects      0
  Input DA rejects          0
  Input SA rejects          0
  Output packet count       0
  Output packet pad count   0
  Output packet error count 0
  CAM destination filters: 2, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Incomplete
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority
Limit
                                %      bps      %      usec
  0 best-effort           95      950000000    95      0      low
none
  3 network-control       5       50000000     5      0      low
none
Interface transmit statistics: Disabled

Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)
Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 0      0 bps
  Output bytes : 0      0 bps
  Input packets: 0      0 pps
  Output packets: 0      0 pps
Security: Zone: public

```

```

Flow Statistics :
Flow Input statistics :
    Self packets :                0
    ICMP packets :                0
    VPN packets :                 0
    Multicast packets :           0
    Bytes permitted by policy :    0
    Connections established :      0
Flow Output statistics:
    Multicast packets :           0
    Bytes permitted by policy :    0
Flow error statistics (Packets dropped due to):
    Address spoofing:             0
    Authentication failed:        0
    Incoming NAT errors:         0
    Invalid zone received packet: 0
    Multiple user authentications: 0
    Multiple incoming NAT:        0
    No parent for a gate:         0
    No one interested in self packets: 0
    No minor session:             0
    No more sessions:             0
    No NAT gate:                  0
    No route present:            0
    No SA for incoming SPI:       0
    No tunnel found:              0
    No session for a gate:        0
    No zone or NULL zone binding  0
    Policy denied:                0
    Security association not active: 0
    TCP sequence number out of window: 0
    Syn-attack protection:        0
    User authentication errors:    0
Protocol inet, MTU: 1500, Generation: 150, Route table: 0
  Flags: Sendbroadcast-pkt-to-re
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
    Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255,
    Generation: 150

```

show interfaces terse

```
user@host> show interfaces terse
```

Interface	Admin	Link	Proto	Local	Remote
ge-0/0/0	up	up			
ge-0/0/0.0	up	up	inet	10.209.4.61/18	
gr-0/0/0	up	up			
ip-0/0/0	up	up			
st0	up	up			
st0.1	up	ready	inet		
ls-0/0/0	up	up			
lt-0/0/0	up	up			
mt-0/0/0	up	up			
pd-0/0/0	up	up			
pe-0/0/0	up	up			
e3-1/0/0	up	up			
t3-2/0/0	up	up			
e1-3/0/0	up	up			
se-4/0/0	up	down			
t1-5/0/0	up	up			
br-6/0/0	up	up			
dc-6/0/0	up	up			
dc-6/0/0.32767	up	up			
bc-6/0/0:1	down	up			
bc-6/0/0:1.0	up	down			
dl0	up	up			
dl0.0	up	up	inet		
dsc	up	up			
gre	up	up			
ipip	up	up			
lo0	up	up			
lo0.16385	up	up	inet	10.0.0.1	--> 0/0
				10.0.0.16	--> 0/0
lsi	up	up			
mtun	up	up			
pimd	up	up			
pime	up	up			
pp0	up	up			

show interfaces terse (vSRX and vSRX 3.0)

user@host> show interfaces terse

Interface	Admin	Link	Proto	Local	Remote
ge-0/0/0	up	up			
ge-0/0/0.0	up	up	inet	1.1.65.1/24	
ge-0/0/1	up	up			


```

ge-0/0/2          up    up
e-0/0/3           up    up
ge-0/0/4          up    up

```

show interfaces controller (Channelized E1 IQ with Logical E1)

```
user@host> show interfaces controller ce1-1/2/6
```

```

Controller                                     Admin Link
ce1-1/2/6                                     up    up
e1-1/2/6                                     up    up

```

show interfaces controller (Channelized E1 IQ with Logical DS0)

```
user@host> show interfaces controller ce1-1/2/3
```

```

Controller                                     Admin Link
ce1-1/2/3                                     up    up
ds-1/2/3:1                                   up    up
ds-1/2/3:2                                   up    up

```

show interfaces descriptions

```
user@host> show interfaces descriptions
```

```

Interface      Admin Link Description
so-1/0/0       up    up    M20-3#1
so-2/0/0       up    up    GSR-12#1
ge-3/0/0       up    up    SMB-OSPF_Area300
so-3/3/0       up    up    GSR-13#1
so-3/3/1       up    up    GSR-13#2
ge-4/0/0       up    up    T320-7#1
ge-5/0/0       up    up    T320-7#2
so-7/1/0       up    up    M160-6#1
ge-8/0/0       up    up    T320-7#3
ge-9/0/0       up    up    T320-7#4
so-10/0/0      up    up    M160-6#2
so-13/0/0      up    up    M20-3#2
so-14/0/0      up    up    GSR-12#2

```

```

ge-15/0/0      up    up    SMB-OSPF_Areal00
ge-15/0/1      up    up    GSR-13#3

```

show interfaces destination-class all

user@host> show interfaces destination-class all

```

Logical interface so-4/0/0.0

      Destination class      Packets      Bytes
                        (packet-per-second)  (bits-per-second)
      gold                  0              0
      (                    0) (              0)
      silver                0              0
      (                    0) (              0)

Logical interface so-0/1/3.0

      Destination class      Packets      Bytes
                        (packet-per-second)  (bits-per-second)
      gold                  0              0
      (                    0) (              0)
      silver                0              0
      (                    0) (              0)

```

show interfaces diagnostics optics

user@host> show interfaces diagnostics optics ge-2/0/0

```

Physical interface: ge-2/0/0
  Laser bias current          : 7.408 mA
  Laser output power          : 0.3500 mW / -4.56 dBm
  Module temperature          : 23 degrees C / 73 degrees F
  Module voltage              : 3.3450 V
  Receiver signal average optical power : 0.0002 mW / -36.99 dBm
  Laser bias current high alarm : Off
  Laser bias current low alarm  : Off
  Laser bias current high warning : Off
  Laser bias current low warning : Off
  Laser output power high alarm : Off
  Laser output power low alarm  : Off
  Laser output power high warning : Off
  Laser output power low warning : Off
  Module temperature high alarm : Off
  Module temperature low alarm  : Off
  Module temperature high warning : Off

```

```

Module temperature low warning      : Off
Module voltage high alarm           : Off
Module voltage low alarm            : Off
Module voltage high warning         : Off
Module voltage low warning          : Off
Laser rx power high alarm           : Off
Laser rx power low alarm            : On
Laser rx power high warning         : Off
Laser rx power low warning          : On
Laser bias current high alarm threshold : 17.000 mA
Laser bias current low alarm threshold : 1.000 mA
Laser bias current high warning threshold : 14.000 mA
Laser bias current low warning threshold : 2.000 mA
Laser output power high alarm threshold : 0.6310 mW / -2.00 dBm
Laser output power low alarm threshold : 0.0670 mW / -11.74 dBm
Laser output power high warning threshold : 0.6310 mW / -2.00 dBm
Laser output power low warning threshold : 0.0790 mW / -11.02 dBm
Module temperature high alarm threshold : 95 degrees C / 203 degrees F
Module temperature low alarm threshold : -25 degrees C / -13 degrees F
Module temperature high warning threshold : 90 degrees C / 194 degrees F
Module temperature low warning threshold : -20 degrees C / -4 degrees F
Module voltage high alarm threshold   : 3.900 V
Module voltage low alarm threshold    : 2.700 V
Module voltage high warning threshold : 3.700 V
Module voltage low warning threshold  : 2.900 V
Laser rx power high alarm threshold   : 1.2590 mW / 1.00 dBm
Laser rx power low alarm threshold    : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold : 0.7940 mW / -1.00 dBm
Laser rx power low warning threshold  : 0.0158 mW / -18.01 dBm

```

show interfaces far-end-interval coc12-5/2/0

```
user@host> show interfaces far-end-interval coc12-5/2/0
```

```

Physical interface: coc12-5/2/0, SNMP ifIndex: 121
05:30-current:
    ES-L: 1, SES-L: 1, UAS-L: 0
05:15-05:30:
    ES-L: 0, SES-L: 0, UAS-L: 0
05:00-05:15:
    ES-L: 0, SES-L: 0, UAS-L: 0
04:45-05:00:
    ES-L: 0, SES-L: 0, UAS-L: 0

```

```

04:30-04:45:
    ES-L: 0, SES-L: 0, UAS-L: 0
04:15-04:30:
    ES-L: 0, SES-L: 0, UAS-L: 0
04:00-04:15:
...

```

show interfaces far-end-interval coc1-5/2/1:1

user@host> run show interfaces far-end-interval coc1-5/2/1:1

```

Physical interface: coc1-5/2/1:1, SNMP ifIndex: 342
05:30-current:
    ES-L: 1, SES-L: 1, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
05:15-05:30:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
05:00-05:15:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
04:45-05:00:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
04:30-04:45:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
04:15-04:30:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
04:00-04:15:

```

show interfaces filters

user@host> show interfaces filters

Interface	Admin	Link	Proto	Input Filter	Output Filter
ge-0/0/0	up	up			
ge-0/0/0.0	up	up	inet		
			iso		
ge-5/0/0	up	up			
ge-5/0/0.0	up	up	any		f-any
			inet		f-inet
			multiservice		
gr-0/3/0	up	up			
ip-0/3/0	up	up			
mt-0/3/0	up	up			
pd-0/3/0	up	up			
pe-0/3/0	up	up			

```

vt-0/3/0      up    up
at-1/0/0      up    up
at-1/0/0.0    up    up    inet
                    iso
at-1/1/0      up    down
at-1/1/0.0    up    down inet
                    iso
....

```

show interfaces flow-statistics (Gigabit Ethernet)

user@host> show interfaces flow-statistics ge-0/0/1.0

```

Logical interface ge-0/0/1.0 (Index 70) (SNMP ifIndex 49)
  Flags: SNMP-Traps Encapsulation: ENET2
  Input packets : 5161
  Output packets: 83
  Security: Zone: zone2
  Allowed host-inbound traffic : bootp bfd bgp  dns dvmrp  ldp msdp nhrp ospf
pgm
  pim rip router-discovery rsvp sap vrrp dhcp finger ftp tftp ident-reset http
https ike
  netconf ping rlogin rpm rsh snmp snmp-trap ssh telnet traceroute xnm-clear-text
xnm-ssl
  lsping
  Flow Statistics :
  Flow Input statistics :
    Self packets :                0
    ICMP packets :                0
    VPN packets :                2564
    Bytes permitted by policy :    3478
    Connections established :      1
  Flow Output statistics:
    Multicast packets :           0
    Bytes permitted by policy :    16994
  Flow error statistics (Packets dropped due to):
    Address spoofing:             0
    Authentication failed:        0
    Incoming NAT errors:          0
    Invalid zone received packet:  0
    Multiple user authentications: 0
    Multiple incoming NAT:         0
    No parent for a gate:          0
    No one interested in self packets: 0

```

```

    No minor session:                0
    No more sessions:                0
    No NAT gate:                     0
    No route present:                0
    No SA for incoming SPI:          0
    No tunnel found:                 0
    No session for a gate:            0
    No zone or NULL zone binding     0
    Policy denied:                   0
    Security association not active:  0
    TCP sequence number out of window: 0
    Syn-attack protection:           0
    User authentication errors:       0
Protocol inet, MTU: 1500
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
    Destination: 203.0.113.1/24, Local: 203.0.113.2, Broadcast: 2.2.2.255

```

show interfaces interval (Channelized OC12)

user@host> **show interfaces interval t3-0/3/0:0**

```

Physical interface: t3-0/3/0:0, SNMP ifIndex: 23
17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
17:28-17:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
17:13-17:28:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
16:58-17:13:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
...
Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
    CES: 230, CSES: 230, SEFS: 230, UAS: 238

```

show interfaces interval (E3)

user@host> **show interfaces interval e3-0/3/0**

```
Physical interface: e3-0/3/0, SNMP ifIndex: 23
 17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
 17:28-17:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
 17:13-17:28:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
 16:58-17:13:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
 16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    ....
Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
    CES: 230, CSES: 230, SEFS: 230, UAS: 238
```

show interfaces interval (SONET/SDH) (SRX devices)

user@host> **show interfaces interval so-0/1/0**

```
Physical interface: so-0/1/0, SNMP ifIndex: 19
 20:02-current:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0,
    SES-P: 0, UAS-P: 0
 19:47-20:02:
    ES-S: 267, SES-S: 267, SEFS-S: 267, ES-L: 267, SES-L: 267, UAS-L: 267,
    ES-P: 267, SES-P: 267, UAS-P: 267
 19:32-19:47:
    ES-S: 56, SES-S: 56, SEFS-S: 56, ES-L: 56, SES-L: 56, UAS-L: 46, ES-P: 56,
    SES-P: 56, UAS-P: 46
 19:17-19:32:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0,
    SES-P: 0, UAS-P: 0
 19:02-19:17:
    ....
```

show interfaces load-balancing (SRX devices)

```
user@host> show interfaces load-balancing
```

Interface	State	Last change	Member count
ams0	Up	1d 00:50	2
ams1	Up	00:00:59	2

show interfaces load-balancing detail (SRX devices)

```
user@host> show interfaces load-balancing detail
```

```
Load-balancing interfaces detail
Interface      : ams0
State          : Up
Last change    : 1d 00:51
Member count   : 2
Members       :
  Interface    Weight  State
  mams-2/0/0   10      Active
  mams-2/1/0   10      Active
```

show interfaces mac-database (All MAC Addresses on a Port SRX devices)

```
user@host> show interfaces mac-database xe-0/3/3
```

```
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
  Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
MAC address      Input frames  Input bytes  Output frames  Output bytes
00:00:00:00:00:00      1           56           0             0
00:00:c0:01:01:02    7023810     323095260    0             0
00:00:c0:01:01:03    7023810     323095260    0             0
00:00:c0:01:01:04    7023810     323095260    0             0
00:00:c0:01:01:05    7023810     323095260    0             0
```


00:00:c0:01:01:06	7023810	323095260	0	0
00:00:c0:01:01:07	7023810	323095260	0	0
00:00:c0:01:01:08	7023809	323095214	0	0
00:00:c0:01:01:09	7023809	323095214	0	0
00:00:c0:01:01:0a	7023809	323095214	0	0
00:00:c0:01:01:0b	7023809	323095214	0	0
00:00:c8:01:01:02	30424784	1399540064	37448598	1722635508
00:00:c8:01:01:03	30424784	1399540064	37448598	1722635508
00:00:c8:01:01:04	30424716	1399536936	37448523	1722632058
00:00:c8:01:01:05	30424789	1399540294	37448598	1722635508
00:00:c8:01:01:06	30424788	1399540248	37448597	1722635462
00:00:c8:01:01:07	30424783	1399540018	37448597	1722635462
00:00:c8:01:01:08	30424783	1399540018	37448596	1722635416
00:00:c8:01:01:09	8836796	406492616	8836795	406492570
00:00:c8:01:01:0a	30424712	1399536752	37448521	1722631966
00:00:c8:01:01:0b	30424715	1399536890	37448523	1722632058

Number of MAC addresses : 21

show interfaces mac-database (All MAC Addresses on a Service SRX devices)

user@host> show interfaces mac-database xe-0/3/3

```
Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
  Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
```

MAC address	Input frames	Input bytes	Output frames	Output bytes
00:00:00:00:00:00	1	56	0	0
00:00:c0:01:01:02	7023810	323095260	0	0
00:00:c0:01:01:03	7023810	323095260	0	0
00:00:c0:01:01:04	7023810	323095260	0	0
00:00:c0:01:01:05	7023810	323095260	0	0
00:00:c0:01:01:06	7023810	323095260	0	0
00:00:c0:01:01:07	7023810	323095260	0	0
00:00:c0:01:01:08	7023809	323095214	0	0
00:00:c0:01:01:09	7023809	323095214	0	0
00:00:c0:01:01:0a	7023809	323095214	0	0
00:00:c0:01:01:0b	7023809	323095214	0	0
00:00:c8:01:01:02	31016568	1426762128	38040381	1749857526
00:00:c8:01:01:03	31016568	1426762128	38040382	1749857572
00:00:c8:01:01:04	31016499	1426758954	38040306	1749854076
00:00:c8:01:01:05	31016573	1426762358	38040381	1749857526
00:00:c8:01:01:06	31016573	1426762358	38040381	1749857526
00:00:c8:01:01:07	31016567	1426762082	38040380	1749857480
00:00:c8:01:01:08	31016567	1426762082	38040379	1749857434

00:00:c8:01:01:09	9428580	433714680	9428580	433714680
00:00:c8:01:01:0a	31016496	1426758816	38040304	1749853984
00:00:c8:01:01:0b	31016498	1426758908	38040307	1749854122

show interfaces mac-database mac-address

user@host> **show interfaces mac-database xe-0/3/3 mac-address (SRX devices) 00:00:c8:01:01:09**

```
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags       : None

Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
  Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address: 00:00:c8:01:01:09, Type: Configured,
    Input bytes      : 202324652
    Output bytes     : 202324560
    Input frames     : 4398362
    Output frames    : 4398360
  Policer statistics:
  Policer type      Discarded frames   Discarded bytes
  Output aggregate      3992386         183649756
```

show interfaces mc-ae (SRX devices)

user@host> **show interfaces mc-ae ae0 unit 512**

```
Member Links      : ae0
Local Status      : active
Peer Status       : active
Logical Interface      : ae0.512
Core Facing Interface : Label Ethernet Interface
ICL-PL            : Label Ethernet Interface
```

show interfaces media (SONET/SDH)

The following example displays the output fields unique to the **show interfaces media** command for a SONET interface (with no level of output specified):

user@host> **show interfaces media so-4/1/2**

```
Physical interface: so-4/1/2, Enabled, Physical link is Up
  Interface index: 168, SNMP ifIndex: 495
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags      : Present Running
  Interface flags: Point-To-Point SNMP-Traps 16384
  Link flags        : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 1783 (00:00:00 ago), Output: 1786 (00:00:08 ago)
  LCP state: Opened
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
  mpls: Not-configured
  CHAP state: Not-configured
  CoS queues        : 8 supported
  Last flapped      : 2005-06-15 12:14:59 PDT (04:31:29 ago)
  Input rate        : 0 bps (0 pps)
  Output rate       : 0 bps (0 pps)
  SONET alarms      : None
  SONET defects     : None
  SONET errors:
    BIP-B1: 121, BIP-B2: 916, REI-L: 0, BIP-B3: 137, REI-P: 16747, BIP-BIP2: 0
  Received path trace: routerb so-1/1/2
  Transmitted path trace: routera so-4/1/2
```

show interfaces policers (SRX devices)

user@host> **show interfaces policers**

Interface	Admin	Link	Proto	Input	Policer	Output	Policer
ge-0/0/0	up	up					
ge-0/0/0.0	up	up	inet				
			iso				
gr-0/3/0	up	up					
ip-0/3/0	up	up					
mt-0/3/0	up	up					
pd-0/3/0	up	up					
pe-0/3/0	up	up					
...							
so-2/0/0	up	up					
so-2/0/0.0	up	up	inet	so-2/0/0.0-in-policer		so-2/0/0.0-out-policer	
			iso				

```
so-2/1/0      up    down
...
```

show interfaces policers interface-name (SRX devices)

```
user@host> show interfaces policers so-2/1/0
```

Interface	Admin	Link	Proto	Input Policer	Output Policer
so-2/1/0	up	down			
so-2/1/0.0	up	down	inet	so-2/1/0.0-in-policer	so-2/1/0.0-out-policer
			iso		
			inet6		

show interfaces queue (SRX devices)

The following truncated example shows the CoS queue sizes for queues 0, 1, and 3. Queue 1 has a queue buffer size (guaranteed allocated memory) of 9192 bytes.

```
user@host> show interfaces queue
```

```
Physical interface: ge-0/0/0, Enabled, Physical link is Up
  Interface index: 134, SNMP ifIndex: 509
Forwarding classes: 8 supported, 8 in use
Egress queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: class0
  Queued:
    Packets          :                0          0 pps
    Bytes            :                0          0 bps
  Transmitted:
    Packets          :                0          0 pps
    Bytes            :                0          0 bps
    Tail-dropped packets :                0          0 pps
    RL-dropped packets  :                0          0 pps
    RL-dropped bytes    :                0          0 bps
    RED-dropped packets :                0          0 pps
    Low                :                0          0 pps
    Medium-low         :                0          0 pps
    Medium-high        :                0          0 pps
    High               :                0          0 pps
    RED-dropped bytes   :                0          0 bps
    Low                :                0          0 bps
    Medium-low         :                0          0 bps
```

```

    Medium-high      :                0                0 bps
    High             :                0                0 bps
Queue Buffer Usage:
    Reserved buffer   :            118750000 bytes
    Queue-depth bytes :
    Current           :                0
..
..
Queue: 1, Forwarding classes: class1
..
..
Queue Buffer Usage:
    Reserved buffer   :                9192 bytes
    Queue-depth bytes :
    Current           :                0
..
..
Queue: 3, Forwarding classes: class3
    Queued:
..
..
Queue Buffer Usage:
    Reserved buffer   :            6250000 bytes
    Queue-depth bytes :
    Current           :                0
..
..

```

show interfaces redundancy (SRX devices)

```
user@host> show interfaces redundancy
```

Interface	State	Last change	Primary	Secondary	Current status
rsp0	Not present		sp-1/0/0	sp-0/2/0	both down
rsp1	On secondary	1d 23:56	sp-1/2/0	sp-0/3/0	primary down
rsp2	On primary	10:10:27	sp-1/3/0	sp-0/2/0	secondary down
rlsq0	On primary	00:06:24	lsq-0/3/0	lsq-1/0/0	both up

show interfaces redundancy (Aggregated Ethernet SRX devices)

```
user@host> show interfaces redundancy
```

Interface	State	Last change	Primary	Secondary	Current status
rlsq0	On secondary	00:56:12	lsq-4/0/0	lsq-3/0/0	both up
ae0					
ae1					
ae2					
ae3					
ae4					

show interfaces redundancy detail (SRX devices)

user@host> show interfaces redundancy detail

```

Interface      : rlsq0
  State        : On primary
  Last change   : 00:45:47
  Primary       : lsq-0/2/0
  Secondary     : lsq-1/2/0
  Current status : both up
  Mode          : hot-standby

Interface      : rlsq0:0
  State        : On primary
  Last change   : 00:45:46
  Primary       : lsq-0/2/0:0
  Secondary     : lsq-1/2/0:0
  Current status : both up
  Mode          : warm-standby

```

show interfaces routing brief (SRX devices)

user@host> show interfaces routing brief

Interface	State	Addresses
so-5/0/3.0	Down	ISO enabled
so-5/0/2.0	Up	MPLS enabled
		ISO enabled
		INET 192.168.2.120
		INET enabled
so-5/0/1.0	Up	MPLS enabled
		ISO enabled
		INET 192.168.2.130
		INET enabled
at-1/0/0.3	Up	CCC enabled

```

at-1/0/0.2      Up    CCC    enabled
at-1/0/0.0      Up    ISO    enabled
                INET   192.168.90.10
                INET   enabled
lo0.0           Up    ISO    47.0005.80ff.f800.0000.0108.0001.1921.6800.5061.00
                ISO    enabled
                INET   127.0.0.1
fxp1.0          Up
fxp0.0          Up    INET   192.168.6.90

```

show interfaces routing detail (SRX devices)

user@host> show interfaces routing detail

```

so-5/0/3.0
  Index: 15, Refcount: 2, State: Up <Broadcast PointToPoint Multicast> Change:<>
  Metric: 0, Up/down transitions: 0, Full-duplex
  Link layer: HDLC serial line Encapsulation: PPP Bandwidth: 155Mbps
  ISO address (null)
    State: <Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
so-5/0/2.0
  Index: 14, Refcount: 7, State: <Up Broadcast PointToPoint Multicast> Change:<>
  Metric: 0, Up/down transitions: 0, Full-duplex
  Link layer: HDLC serial line Encapsulation: PPP Bandwidth: 155Mbps
  MPLS address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4458 bytes
  ISO address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
  INET address 192.168.2.120
    State: <Up Broadcast PointToPoint Multicast Localup> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
    Local address: 192.168.2.120
    Destination: 192.168.2.110/32
  INET address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
...

```

show interfaces routing-instance all (SRX devices)

user@host> show interfaces terse routing-instance all

Interface	Admin	Link	Proto	Local	Remote	Instance
at-0/0/1	up	up	inet	10.0.0.1/24		
ge-0/0/0.0	up	up	inet	192.168.4.28/24		sample-a
at-0/1/0.0	up	up	inet6	fe80::a:0:0:4/64		sample-b
so-0/0/0.0	up	up	inet	10.0.0.1/32		

show interfaces snmp-index (SRX devices)

user@host> show interfaces snmp-index 33

```
Physical interface: so-2/1/1, Enabled, Physical link is Down
Interface index: 149, SNMP ifIndex: 33
Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,
Loopback: None, FCS: 16, Payload scrambler: Enabled
Device flags   : Present Running Down
Interface flags: Hardware-Down Point-To-Point SNMP-Traps 16384
Link flags     : Keepalives
CoS queues     : 8 supported
Last flapped   : 2005-06-15 11:45:57 PDT (05:38:43 ago)
Input rate     : 0 bps (0 pps)
Output rate    : 0 bps (0 pps)
SONET alarms   : LOL, PLL, LOS
SONET defects  : LOL, PLL, LOF, LOS, SEF, AIS-L, AIS-P
```

show interfaces source-class all (SRX devices)

user@host> show interfaces source-class all

```
Logical interface so-0/1/0.0

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

Logical interface so-0/1/3.0

          Packets          Bytes
Source class      (packet-per-second) (bits-per-second)
          gold              0              0
                  (            0) (            0)
          bronze              0              0
```


	(0)	(0)
silver		116113		9753492
	(939)	(631616)

show interfaces statistics (Fast Ethernet SRX devices)

user@host> show interfaces fe-1/3/1 statistics

```
Physical interface: fe-1/3/1, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 1042
  Description: ford fe-1/3/1
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:90:69:93:04:dc, Hardware address: 00:90:69:93:04:dc
  Last flapped   : 2006-04-18 03:08:59 PDT (00:01:24 ago)
  Statistics last cleared: Never
  Input rate      : 0 bps (0 pps)
  Output rate     : 0 bps (0 pps)
  Input errors: 0, Output errors: 0
  Active alarms   : None
  Active defects  : None
  Logical interface fe-1/3/1.0 (Index 69) (SNMP ifIndex 50)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500
      Flags: Is-Primary, DCU, SCU-in
        Destination class      Packets      Bytes
                                (packet-per-second) (bits-per-second)
        silver1                 0              0
                                (0) (0)
        silver2                 0              0
                                (0) (0)
        silver3                 0              0
                                (0) (0)
    Addresses, Flags: Is-Default Is-Preferred Is-Primary
      Destination: 10.27.245/24, Local: 10.27.245.2,
      Broadcast: 10.27.245.255
    Protocol iso, MTU: 1497
      Flags: Is-Primary
```

Physical interface: et-0/1/0, SNMP ifIndex 515				
14:45-current		Elapse time:900 Seconds		
Near End	Suspect Flag:False		Reason:None	
PM	COUNT	THRESHOLD	TCA-ENABLED	TCA-RAISED
OTU-BBE	0	800	No	No
OTU-ES	0	135	No	No
OTU-SES	0	90	No	No
OTU-UAS	427	90	No	No
Far End	Suspect Flag:True		Reason:Unknown	
PM	COUNT	THRESHOLD	TCA-ENABLED	TCA-RAISED
OTU-BBE	0	800	No	No

OTU-ES	0	135	No	No
OTU-SES	0	90	No	No
OTU-UAS	0	90	No	No
Near End	Suspect Flag:False		Reason:None	
PM	COUNT	THRESHOLD	TCA-ENABLED	TCA-RAISED
ODU-BBE	0	800	No	No
ODU-ES	0	135	No	No
ODU-SES	0	90	No	No
ODU-UAS	427	90	No	No
Far End	Suspect Flag:True		Reason:Unknown	
PM	COUNT	THRESHOLD	TCA-ENABLED	TCA-RAISED
ODU-BBE	0	800	No	No
ODU-ES	0	135	No	No
ODU-SES	0	90	No	No
ODU-UAS	0	90	No	No
FEC	Suspect Flag:False		Reason:None	
PM	COUNT	THRESHOLD	TCA-ENABLED	TCA-RAISED
FEC-CorrectedErr	2008544300	0	NA	NA
FEC-UncorrectedWords	0	0	NA	NA
BER	Suspect Flag:False		Reason:None	
PM	MIN	MAX	AVG	THRESHOLD
TCA-RAISED				
BER	3.6e-5	5.8e-5	3.6e-5	10.0e-3
Yes				No
Physical interface: et-0/1/0, SNMP ifIndex 515				
14:45-current				
Suspect Flag:True	Reason:Object Disabled			
PM	CURRENT	MIN	MAX	AVG
TCA-ENABLED	TCA-RAISED			
				(MIN)
(MAX)	(MIN)	(MAX)	(MIN)	(MAX)
Lane chromatic dispersion	0	0	0	0
0	NA	NA	NA	NA
Lane differential group delay	0	0	0	0
0	NA	NA	NA	NA
q Value	120	120	120	120
0	NA	NA	NA	NA
SNR	28	28	29	28
0	NA	NA	NA	NA
Tx output power(0.01dBm)	-5000	-5000	-5000	-5000
-100	No	No	No	No

Rx input power(0.01dBm)	-3642	-3665	-3626	-3637	-1800
-500 No No No No					
Module temperature(Celsius)	46	46	46	46	-5
75 No No No No					
Tx laser bias current(0.1mA)	0	0	0	0	0
0 NA NA NA NA					
Rx laser bias current(0.1mA)	1270	1270	1270	1270	0
0 NA NA NA NA					
Carrier frequency offset(MHz)	-186	-186	-186	-186	-5000
5000 No No No No					

show security zones (SRX devices)

user@host> show security zones

```

Functional zone: management
  Description: This is the management zone.
  Policy configurable: No
  Interfaces bound: 1
  Interfaces:
    ge-0/0/0.0
Security zone: Host
  Description: This is the host zone.
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
    fxp0.0
Security zone: abc
  Description: This is the abc zone.
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
    ge-0/0/1.0
Security zone: def
  Description: This is the def zone.
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
    ge-0/0/2.0

```

show interfaces (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 757](#)

[show interfaces brief \(Loopback\) on page 758](#)

[show interfaces detail \(Loopback\) on page 759](#)

[show interfaces extensive \(Loopback\) on page 760](#)

Output Fields

Table 30 on page 754 lists the output fields for the **show interfaces** (loopback) command. Output fields are listed in the approximate order in which they appear.

Table 30: Loopback show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical Interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .	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 <i>Common Output Fields Description</i> .	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .	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 <i>Common Output Fields Description</i> .	detail extensive none
Physical info	Information about the physical interface.	detail extensive

Table 30: Loopback show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface. • Input packets, Output packets—Number of packets received and transmitted on the interface. 	detail extensive
Input errors	<ul style="list-style-type: none"> • 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

Table 30: Loopback show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Output errors	<ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly, 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 <i>Common Output Fields Description</i> .	brief detail extensive
Encapsulation	Encapsulation on the logical interface.	brief 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
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

Table 30: Loopback show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
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 <i>Common Output Fields Description</i> .	detail extensive none
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	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: 203.0.113.1
    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  203.0.113.1          --> 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: 203.0.113.1, 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: 203.0.113.1, 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 (PPPoE)

Syntax

```
show interfaces pp0.logical
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Release Information

Command introduced before Junos OS Release 7.4.

Description

(M120 routers, M320 routers, and MX Series routers only). Display status information about the PPPoE interface.

Options

pp0.logical—Display standard status information about the PPPoE interface.

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

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about PPPoE interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display PPPoE interface statistics.

Required Privilege Level

view

List of Sample Output

[show interfaces \(PPPoE\) on page 771](#)

[show interfaces \(PPPoE over Aggregated Ethernet\) on page 771](#)

[show interfaces brief \(PPPoE\) on page 772](#)

[show interfaces detail \(PPPoE\) on page 772](#)

[show interfaces extensive \(PPPoE on M120 and M320 Routers\) on page 773](#)

Output Fields

[Table 31 on page 764](#) lists the output fields for the **show interfaces (PPPoE)** command. Output fields are listed in the approximate order in which they appear.

Table 31: 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 <i>Common Output Fields Description</i> .	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	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 <i>Common Output Fields Description</i> .	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .	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 <i>Common Output Fields Description</i> .	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

Table 31: show interfaces (PPPoE) Output Fields (*continued*)

Field Name	Field Description	Level of Output
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
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive
IPv6 transit statistics	<p>Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled.</p> <p>NOTE: These fields include dropped traffic and exception traffic, as those fields are not separately defined.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive

Table 31: show interfaces (PPPoE) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Input errors	<p>Input errors on the interface:</p> <ul style="list-style-type: none"> • Errors—Sum of incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Giants—Number of frames received that are larger than the giant threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • Resource errors—Sum of B chip Tx drops and IXP Tx net transmit drops. 	extensive
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions —Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and 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		
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

Table 31: show interfaces (PPPoE) Output Fields (*continued*)

Field Name	Field Description	Level of Output
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 <i>Common Output Fields Description</i> .	All levels
Encapsulation	Type of encapsulation configured on the logical interface.	All levels
PPP parameters	PPP status: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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 31: show interfaces (PPPoE) Output Fields (*continued*)

Field Name	Field Description	Level of Output
IPv6 transit statistics	<p>Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.</p> <p>NOTE: The packet and byte counts in these fields include traffic that is dropped and does not leave the router.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive
Local statistics	<p>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.</p>	detail extensive
Transit statistics	<p>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.</p> <p>NOTE: The packet and byte counts in these fields include traffic that is dropped and does not leave the router.</p>	detail extensive
Keepalive settings	<p>(PPP and HDLC) Configured settings for keepalives.</p> <ul style="list-style-type: none"> • 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-countnumber—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

Table 31: show interfaces (PPPoE) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Keepalive statistics	<p>(PPP and HDLC) Information about keepalive packets.</p> <ul style="list-style-type: none"> • Input—Number of keepalive packets received by PPP. <ul style="list-style-type: none"> • (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. <ul style="list-style-type: none"> • (last seen 00:00:00 ago)—Time the last keepalive packet was sent, in the format hh:mm:ss. <p>(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.</p>	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
LCP state	<p>(PPP) Link Control Protocol state.</p> <ul style="list-style-type: none"> • Conf-ack-received—Acknowledgement was received. • Conf-ack-sent—Acknowledgement was sent. • Conf-req-sent—Request was sent. • Down—LCP negotiation is incomplete (not yet completed or has failed). • Not-configured—LCP is not configured on the interface. • Opened—LCP negotiation is successful. 	none detail extensive
NCP state	<p>(PPP) Network Control Protocol state.</p> <ul style="list-style-type: none"> • Conf-ack-received—Acknowledgement was received. • Conf-ack-sent—Acknowledgement was sent. • Conf-req-sent—Request was sent. • Down—NCP negotiation is incomplete (not yet completed or has failed). • Not-configured—NCP is not configured on the interface. • Opened—NCP negotiation is successful. 	detail extensive none

Table 31: show interfaces (PPPoE) Output Fields (*continued*)

Field Name	Field Description	Level of Output
CHAP state	<p>(PPP) Displays the state of the Challenge Handshake Authentication Protocol (CHAP) during its transaction.</p> <ul style="list-style-type: none"> • Chap-Chal-received—Challenge was received but response not yet sent. • Chap-Chal-sent—Challenge was sent. • Chap-Resp-received—Response was received for the challenge sent, but CHAP has not yet moved into the Success state. (Most likely with RADIUS authentication.) • Chap-Resp-sent—Response was sent for the challenge received. • Closed—CHAP authentication is incomplete. • Failure—CHAP authentication failed. • Not-configured—CHAP is not configured on the interface. • Success—CHAP authentication was successful. 	none detail extensive
Protocol	Protocol family configured on the logical interface.	detail extensive none
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
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 <i>Common Output Fields Description</i> .	detail extensive none
Addresses, Flags	Information about the addresses configured for the protocol family. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address.	detail extensive none

Sample Output

show interfaces (PPPoE)

user@host> show interfaces pp0

```
Physical interface: pp0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 24
  Type: PPPoE, Link-level type: PPPoE, MTU: 1532
  Device flags      : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type        : Full-Duplex
  Link flags       : None
  Input rate       : 0 bps (0 pps)
  Output rate      : 0 bps (0 pps)

Logical interface pp0.0 (Index 72) (SNMP ifIndex 72)
  Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
  PPPoE:
    State: SessionDown, Session ID: None,
    Service name: None, Configured AC name: sapphire,
    Auto-reconnect timeout: 100 seconds, Idle timeout: Never,
    Underlying interface: at-5/0/0.0 (Index 70)
  Input packets : 0
  Output packets: 0
  LCP state: Not-configured
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
  mpls: Not-configured
  CHAP state: Closed
    Protocol inet, MTU: 100
    Flags: User-MTU, Negotiate-Address
```

show interfaces (PPPoE over Aggregated Ethernet)

user@host> show interfaces pp0.1073773821

```
Logical interface pp0.1073773821 (Index 80) (SNMP ifIndex 32584)
  Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
  PPPoE:
    State: SessionUp, Session ID: 1,
    Session AC name: alcor, Remote MAC address: 00:00:5e:00:53:01,
    Underlying interface: demux0.100 (Index 88)
  Link:
    ge-1/0/0.32767
```

```

    ge-1/0/1.32767
    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: Sendbcst-pkt-to-re
    Addresses, Flags: Is-Primary
    Local: 203.0.113.1

```

show interfaces brief (PPPoE)

user@host> **show interfaces pp0 brief**

```

Physical interface: pp0, Enabled, Physical link is Up
  Type: PPPoE, Link-level type: PPPoE, MTU: 1532, Speed: Unspecified
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps

Logical interface pp0.0
  Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
  PPPoE:
    State: SessionDown, Session ID: None,
    Service name: None, Configured AC name: sapphire,
    Auto-reconnect timeout: 100 seconds, Idle timeout: Never,
    Underlying interface: at-5/0/0.0 (Index 70)
  inet

```

show interfaces detail (PPPoE)

user@host> **show interfaces pp0 detail**

```

Physical interface: pp0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 24, Generation: 9
  Type: PPPoE, Link-level type: PPPoE, MTU: 1532, Speed: Unspecified
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified

```



```

Hold-times      : Up 0 ms, Down 0 ms
Current address: Unspecified, Hardware address: Unspecified
Alternate link address: Unspecified
Statistics last cleared: Never
Traffic statistics:
  Input  bytes :                0                0 bps
  Output bytes :                0                0 bps
  Input  packets:                0                0 pps
  Output packets:                0                0 pps
Logical interface pp0.0 (Index 72) (SNMP ifIndex 72) (Generation 14)
  Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
  PPPoE:
    State: SessionDown, Session ID: None,
    Service name: None, Configured AC name: sapphire,
    Auto-reconnect timeout: 100 seconds, Idle timeout: Never,
    Underlying interface: at-5/0/0.0 (Index 70)
  Traffic statistics:
    Input  bytes :                0
    Output bytes :                0
    Input  packets:                0
    Output packets:                0
  Local statistics:
    Input  bytes :                0
    Output bytes :                0
    Input  packets:                0
    Output packets:                0
  Transit statistics:
    Input  bytes :                0                0 bps
    Output bytes :                0                0 bps
    Input  packets:                0                0 pps
    Output packets:                0                0 pps
LCP state: Not-configured
NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
mpls: Not-configured
CHAP state: Closed
  Protocol inet, MTU: 100, Generation: 14, Route table: 0
  Flags: User-MTU, Negotiate-Address

```

show interfaces extensive (PPPoE on M120 and M320 Routers)

user@host> show interfaces pp0 extensive

```

Physical interface: pp0, Enabled, Physical link is Up
Interface index: 128, SNMP ifIndex: 93, Generation: 129

```

```

Type: PPPoE, Link-level type: PPPoE, MTU: 1532, Speed: Unspecified
Device flags   : Present Running
Interface flags: Point-To-Point SNMP-Traps
Link type      : Full-Duplex
Link flags     : None
Physical info  : Unspecified
Hold-times    : Up 0 ms, Down 0 ms
Current address: Unspecified, Hardware address: Unspecified
Alternate link address: Unspecified
Statistics last cleared: Never
Traffic statistics:
  Input  bytes :          972192          0 bps
  Output bytes :          975010          0 bps
  Input  packets:           1338          0 pps
  Output packets:           1473          0 pps
IPv6 transit statistics:
  Input  bytes :              0
  Output bytes :              0
  Input  packets:              0
  Output packets:              0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
0,
  Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0

Logical interface pp0.0 (Index 69) (SNMP ifIndex 96) (Generation 194)
  Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
  PPPoE:
    State: SessionUp, Session ID: 26,
    Session AC name: None, AC MAC address: 00:00:5e:00:53:12,
    Service name: None, Configured AC name: None,
    Auto-reconnect timeout: Never, Idle timeout: Never,
    Underlying interface: ge-3/0/1.0 (Index 67)
  Traffic statistics:
    Input  bytes :          252
    Output bytes :          296
    Input  packets:           7
    Output packets:           8
  IPv6 transit statistics:
    Input  bytes :              0
    Output bytes :              0

```

```

    Input  packets:                0
    Output packets:                0
Local statistics:
    Input  bytes   :                252
    Output bytes   :                296
    Input  packets:                7
    Output packets:                8
Transit statistics:
    Input  bytes   :                0          0 bps
    Output bytes   :                0          0 bps
    Input  packets:                0          0 pps
    Output packets:                0          0 pps
IPv6 transit statistics:
    Input  bytes   :                0
    Output bytes   :                0
    Input  packets:                0
    Output packets:                0
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive statistics:
    Input : 1 (last seen 00:00:00 ago)
    Output: 1 (last sent 00:00:03 ago)
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
CHAP state: Closed
PAP state: Closed
Protocol inet, MTU: 1492, Generation: 171, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
    Destination: 203.0.113.2, Local: 203.0.113.1, Broadcast: Unspecified,
Generation: 206

```

show interfaces demux0 (Demux Interfaces)

Syntax

```
show interfaces demux0.logical-interface-number
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Release Information

Command introduced in Junos OS Release 9.0.

Description

(MX Series and M Series routers only) Display status information about the specified demux interface.

Options

none—Display standard information about the specified demux interface.

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

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index snmp-index—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

RELATED DOCUMENTATION

Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration

List of Sample Output

[show interfaces demux0 \(Demux\) on page 784](#)

[show interfaces demux0 \(PPPoE over Aggregated Ethernet\) on page 785](#)

[show interfaces demux0 extensive \(Targeted Distribution for Aggregated Ethernet Links\) on page 786](#)

[show interfaces demux0 \(ACI Interface Set Configured\) on page 786](#)

Output Fields

Table 32 on page 777 lists the output fields for the **show interfaces demux0** (Demux Interfaces) command. Output fields are listed in the approximate order in which they appear.

Table 32: show interfaces demux0 (Demux 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
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .	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

Table 32: show interfaces demux0 (Demux Interfaces) Output Fields (*continued*)

Field Name	Field Description	Level of Output
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 <i>Common Output Fields Description</i> .	brief detail extensive none
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .	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 <i>Common Output Fields Description</i> .	detail extensive none
Physical info	Information about the physical interface.	detail extensive
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive
Current address	Configured MAC address.	detail extensive
Hardware address	Hardware MAC address.	detail extensive
Alternate link address	Backup address of the link.	detail extensive
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive

Table 32: show interfaces demux0 (Demux Interfaces) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—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. <p>NOTE: These fields include dropped traffic and exception traffic, as those fields are not separately defined.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive
Input errors	<p>Input errors on the interface whose definitions are as follows:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Giants—Number of frames received that are larger than the giant 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 32: show interfaces demux0 (Demux Interfaces) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • 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 <i>Common Output Fields Description</i> .	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

Table 32: show interfaces demux0 (Demux Interfaces) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Demux	<p>Specific IP demultiplexing (demux) values:</p> <ul style="list-style-type: none"> • 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
<i>protocol-family</i>	Protocol family configured on the logical interface.	brief
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface set. • 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. <p>NOTE: The packet and byte counts in these fields include traffic that is dropped and does not leave the router.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive
Local statistics	<p>Number of transit bytes and packets received and transmitted on the local interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive

Table 32: show interfaces demux0 (Demux Interfaces) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Transit statistics	<p>Number and rate of bytes and packets transiting the switch.</p> <p>NOTE: The packet and byte counts in these fields include traffic that is dropped and does not leave the router.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive
IPv6 Transit statistics	<p>Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.</p> <p>NOTE: The packet and byte counts in these fields include traffic that is dropped and does not leave the router.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive
Input packets	Number of packets received on the interface.	none
Output packets	Number of packets transmitted on the interface.	none
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under <i>Common Output Fields Description</i> .	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 <i>Common Output Fields Description</i> .	detail extensive none

Table 32: show interfaces demux0 (Demux Interfaces) Output Fields (*continued*)

Field Name	Field Description	Level of Output
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 <i>Common Output Fields Description</i> .	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
Duplicate Protection	State of duplicate protection: On or Off . Duplicate protection prevents the activation of another dynamic PPPoE logical interface on the same underlying interface when a dynamic PPPoE logical interface for a client with the same MAC address is already active on that interface.	detail extensive none
Direct Connect	State of the configuration to ignore DSL Forum VSAs: On or Off . When configured, the router ignores any of these VSAs received from a directly connected CPE device on the interface.	detail extensive none

Table 32: show interfaces demux0 (Demux Interfaces) Output Fields (*continued*)

Field Name	Field Description	Level of Output
AC Name	Name of the access concentrator.	detail extensive none

Sample Output

show interfaces demux0 (Demux)

user@host> show interfaces demux0

```
Physical interface: demux0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 79, Generation: 129
  Type: Software-Pseudo, Link-level type: Unspecified, MTU: 9192, Clocking: 1,
  Speed: Unspecified
  Device flags      : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type        : Full-Duplex
  Link flags       : None
  Physical info    : Unspecified
  Hold-times       : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped     : Never
  Statistics last cleared: Never
  Traffic statistics:
    Input  bytes   :                0          0 bps
    Output bytes   :                0          0 bps
    Input  packets :                0          0 pps
    Output packets :                0          0 pps
  IPv6 transit statistics:
    Input  bytes   :                0
    Output bytes   :                0
    Input  packets :                0
    Output packets :                0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0
```

```

Logical interface demux0.0 (Index 87) (SNMP ifIndex 84) (Generation 312)
  Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
  Demux:
    Underlying interface: ge-2/0/1.0 (Index 74)
  Family Inet Source prefixes, total 1
  Prefix: 203.0.113/24
  Traffic statistics:
    Input  bytes   :                0
    Output bytes   :             1554
    Input  packets :                0
    Output packets :             37
  IPv6 transit statistics:
    Input  bytes   :                0
    Output bytes   :                0
    Input  packets :                0
    Output packets :                0
  Local statistics:
    Input  bytes   :                0
    Output bytes   :             1554
    Input  packets :                0
    Output packets :             37
  Transit statistics:
    Input  bytes   :                0                0 bps
    Output bytes   :                0                0 bps
    Input  packets :                0                0 pps
    Output packets :                0                0 pps
  IPv6 transit statistics:
    Input  bytes   :                0
    Output bytes   :                0
    Input  packets :                0
    Output packets :                0
  Protocol inet, MTU: 1500, Generation: 395, Route table: 0
  Flags: Is-Primary, Mac-Validate-Strict
  Mac-Validate Failures: Packets: 0, Bytes: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 203.0.113/24, Local: 203.0.113.13, Broadcast: 203.0.113.255,

    Generation: 434

```

show interfaces demux0 (PPPoE over Aggregated Ethernet)

```
user@host> show interfaces demux0.100
```

```

Logical interface demux0.100 (Index 76) (SNMP ifIndex 61160)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ]
  Encapsulation: ENET2
  Demux:
    Underlying interface: ae0 (Index 199)
  Link:
    ge-1/0/0
    ge-1/1/0
  Input packets : 0
  Output packets: 0
  Protocol pppoe
    Dynamic Profile: pppoe-profile,
    Service Name Table: service-table1,
    Max Sessions: 100, Duplicate Protection: On,
    Direct Connect: Off,
    AC Name: pppoe-server-1

```

show interfaces demux0 extensive (Targeted Distribution for Aggregated Ethernet Links)

user@host> **show interfaces demux0.1073741824 extensive**

```

Logical interface demux0.1073741824 (Index 75) (SNMP ifIndex 558) (Generation
346)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ] Encapsulation: ENET2
  Demux:
    Underlying interface: ae0 (Index 201)
  Link:
    ge-1/0/0
    ge-1/1/0
    ge-2/0/7
    ge-2/0/8
  Targeting summary:
    ge-1/1/0, primary, Physical link is Up
    ge-2/0/8, backup, Physical link is Up
  Bandwidth: 1000mbps

```

show interfaces demux0 (ACI Interface Set Configured)

user@host> **show interfaces demux0.1073741827**

```

Logical interface demux0.1073741827 (Index 346) (SNMP ifIndex 527)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1802 0x8100.302 ] Encapsulation:

```

ENET2

Demux: Source Family Inet

ACI VLAN:

Dynamic Profile: aci-vlan-set-profile

Demux:

Underlying interface: ge-1/0/0 (Index 138)

Input packets : 18

Output packets: 16

Protocol inet, MTU: 1500

Flags: Sendbcast-pkt-to-re, Unnumbered

Donor interface: lo0.0 (Index 322)

Preferred source address: 203.0.113.202

Addresses, Flags: Primary Is-Default Is-Primary

Local: 203.0.113.119

Protocol pppoe

Dynamic Profile: aci-vlan-pppoe-profile,

Service Name Table: None,

Max Sessions: 32000, Max Sessions VSA Ignore: Off,

Duplicate Protection: On, Short Cycle Protection: Off,

Direct Connect: Off,

AC Name: nbc

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 Policies, Firewall Filters, and Traffic Policers User Guide*. For related operational mode commands, see the [CLI Explorer](#).

Required Privilege Level

view

List of Sample Output

- [show interfaces filters on page 789](#)
- [show interfaces filters \(Interface-Name\) on page 790](#)
- [show interfaces filters \(PTX Series Packet Transport Routers\) on page 790](#)

Output Fields

[Table 33 on page 788](#) lists the output fields for the **show interfaces filters** command. Output fields are listed in the approximate order in which they appear.

Table 33: 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 .

Table 33: show interfaces filters Output Fields (*continued*)

Field Name	Field Description
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
               up   up   iso
ge-5/0/0       up   up
ge-5/0/0.0     up   up   any           f-any
               up   up   inet          f-inet
               up   up   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
               up   up   iso
at-1/1/0       up   down
at-1/1/0.0     up   down inet
               up   down 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 l2-routing-instance

Syntax

```
show interfaces l2-routing-instance routing-instance-name
```

Release Information

Command introduced in Junos OS Release 16.1R4.

Description

Display information about core-facing physical interfaces in the specified routing instance. The routing instance must be configured for Layer 2 wholesale operations or the command displays no output.

Options

routing-instance-name—Name of the routing instance.

Required Privilege Level

view

RELATED DOCUMENTATION

| [Layer 2 Wholesale with ANCP-Triggered VLANs Overview](#) | 116

List of Sample Output

[show interfaces l2-routing-instance on page 792](#)

Output Fields

[Table 34 on page 791](#) lists the output fields for the **show interfaces l2-routing-instance** command. Output fields are listed in the approximate order in which they appear.

Table 34: show interfaces l2-routing-instance Output Fields

Field Name	Field Description
Interface	Name of .the core-facing physical interface.
VLAN Id	Trunk VLAN ID assigned to the core-facing physical interface within its input VLAN map.
Inner VLAN Id total	Aggregate number of VLAN IDs assigned to this physical interface based on the inner VLAN ID-swap-ranges configured for the interface.

Table 34: show interfaces l2-routing-instance Output Fields (*continued*)

Field Name	Field Description
Inner VLAN range	Inner VLAN ID swap range; numbers that can be used for swapping inner VLAN IDs.
Inner VLAN range total	Size of the inner VLAN ID swap range.
Inner VLAN id use count	Number of inner VLAN ID tags in use.
Inner VLAN id free count	Number of inner VLAN ID tags not in use.

Sample Output

show interfaces l2-routing-instance

user@host> **show interfaces l2-routing-instance NSP1**

```
Interface: ge-1/1/1.0
VLAN Id: 100
  Inner VLAN Id total: 6
  Inner VLAN range: 15-20
    Inner VLAN range total   : 6
    Inner VLAN id use count  : 1
    Inner VLAN id free count : 5
```

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 related operational mode commands for routing instances and protocols, see the [CLI Explorer](#).

Required Privilege Level

view

List of Sample Output

[show interfaces routing brief on page 795](#)

[show interfaces routing brief \(TX Matrix Plus Router\) on page 796](#)

[show interfaces routing detail on page 797](#)

[show interfaces routing detail \(TX Matrix Plus Router\) on page 797](#)

Output Fields

[Table 35 on page 794](#) lists the output fields for the **show interfaces routing** command. Output fields are listed in the approximate order in which they appear.

Table 35: 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
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: <ul style="list-style-type: none"> • 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

Table 35: show interfaces routing Output Fields (*continued*)

Field Name	Field Description	Level of Output
Protocol address	<p>Information about the configuration of protocols on the interface:</p> <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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   203.0.113.91
                ISO    enabled
                MPLS   enabled
ge-23/0/3.0    Up    INET   203.0.113.81
                ISO    enabled
                MPLS   enabled
ge-23/0/2.0    Up    INET   203.0.113.71
                ISO    enabled
                MPLS   enabled
ge-23/0/1.0    Up    INET   203.0.113.61
                ISO    enabled
                MPLS   enabled
ge-23/0/0.0    Up    INET   203.0.113.51
                ISO    enabled
                MPLS   enabled
ge-31/0/7.599  Up    INET   192.0.2.93
ge-31/0/7.598  Up    INET   192.0.2.89
ge-31/0/7.597  Up    INET   192.0.2.85
ge-31/0/7.596  Up    INET   192.0.2.81
ge-31/0/7.595  Up    INET   192.0.2.77
ge-31/0/7.594  Up    INET   192.0.2.73
...
ixgbe1.0       Up    INET   203.0.113.34
                INET   198.51.100.4
                INET6  fe80::200:1ff:fe22:4
                INET6  fec0::a:22:0:4
ixgbe0.0       Up    INET   203.0.113.34
                INET   198.51.100

```



```

                                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 203.0.113.91
  State: <Up Broadcast Multicast Localup> Change: <> Flags: <RT-Change>
  Preference 0, metric 0, MTU 1500 bytes
  Broadcast address 203.0.113.93
  Destination: 203.0.113.0/30
  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 203.0.113.81
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <RT-Change>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 10.8.1.3
    Destination: 203.0.113.80/30
    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 203.0.113.71
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <RT-Change>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 203.0.113.73
    Destination: 203.0.113.70/30

```

```

    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 203.0.113.61
        State: <Up Broadcast Multicast Localup> Change: <> Flags: <RT-Change>
        Preference 0, metric 0, MTU 1500 bytes
        Broadcast address 203.0.113.63
    ...
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 203.0.113.34
        State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
        Preference 0, metric 0, MTU 1500 bytes
        Broadcast address 203.0.113.255
        Destination: 203.0.113.0/8
        System flags: <Is-Preferred>
    INET address 198.51.100.4
        State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
        Preference 0, metric 0, MTU 1500 bytes
        Broadcast address 191.255.255.255
        Destination: 192.0.2.0/2
        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 203.0.113.34
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 203.0.113.255
    Destination: 203.0.113.0/8
    System flags: <Is-Preferred>
  INET address 198.51.100.4
    State: <Up Broadcast Multicast Localup> Change: <> Flags: <>
    Preference 0, metric 0, MTU 1500 bytes
    Broadcast address 191.255.255.255
    Destination: 192.0.2.0/2
    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 interfaces routing-instance

Syntax

```
show interfaces routing-instance (instance-name | all)
<brief | detail | extensive | terse>
```

Release Information

Command introduced in Junos OS Release 9.1.

Description

Display information about the interfaces configured for either a specific routing instance or for all of the routing instances.

Options

all—Display information about all of the interfaces configured for all of the routing instances on the router.

instance-name—Display information about the interfaces configured for the specified routing instance.

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

Required Privilege Level

view

List of Sample Output

[show interfaces routing-instance terse on page 801](#)

[show interfaces routing-instance all on page 802](#)

[show interfaces routing-instance extensive on page 802](#)

Output Fields

The output fields from the **show interfaces routing-instance** command are identical to those produced by the **show interfaces *interface-name*** command. For a description of output fields, see the other chapters in this manual.

Sample Output

show interfaces routing-instance terse

```
user@host> show interfaces routing-instance sample terse
```

Interface	Admin	Link	Proto	Local	Remote
ge-0/0/0.0	up	up	inet	192.168.4.28/24	

Sample Output

show interfaces routing-instance all

user@host> **show interfaces terse routing-instance all**

Interface	Admin	Link	Proto	Local	Remote	Instance
at-0/0/1	up	up	inet	203.0.113.1/24		
ge-0/0/0.0	up	up	inet	192.168.4.28/24		sample-a
at-0/1/0.0	up	up	inet6	fe80::a:0:0:4/64		sample-b
so-0/0/0.0	up	up	inet	203.0.113.1/32		

show interfaces routing-instance extensive

user@host**show interfaces fe-0/1/3 routing-instance instance2 extensive**

```

Logical interface fe-0/1/3.0 (Index 70) (SNMP ifIndex 53) (Generation 211)
  Flags: SNMP-Traps Encapsulation: ENET2
  Traffic statistics:
    Input bytes : 0
    Output bytes : 42
    Input packets: 0
    Output packets: 1
  IPv6 transit statistics:
    Input bytes : 0
    Output bytes : 0
    Input packets: 0
    Output packets: 0
  Local statistics:
    Input bytes : 0
    Output bytes : 42
    Input packets: 0
    Output packets: 1
  Transit statistics:
    Input bytes : 0 0 bps
    Output bytes : 0 0 bps
    Input packets: 0 0 pps
    Output packets: 0 0 pps
  IPv6 transit statistics:
    Input bytes : 0
    Output bytes : 0
    Input packets: 0
    Output packets: 0
  Protocol inet, MTU: 1500, Generation: 252, Route table: 4

```

```
Flags: Is-Primary  
Addresses, Flags: Is-Default Is-Preferred Is-Primary  
Destination: 192.0.2/24, Local: 192.0.2.51, Broadcast: 192.0.2.255,  
Generation: 263
```

show network-access aaa statistics

Syntax

```
show network-access aaa statistics
<accounting (detail)>
<address-assignment (client | pool pool-name)>
<dynamic-requests>
<radius>
<session-limit-per-username>
```

Release Information

Command introduced in Junos OS Release 9.1.

address-assignment option introduced in Junos OS Release 10.0.

radius option introduced in Junos OS Release 11.4.

detail option introduced in Junos OS Release 13.3.

session-limit-per-username option introduced in Junos OS Release 18.4R1 on MX Series routers.

Description

Display AAA accounting, address-assignment, dynamic request statistics, RADIUS settings and statistics, and subscriber session limit statistics.

Options

accounting (detail)—(Optional) Display AAA accounting statistics. The **detail** keyword displays additional accounting information

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.

session-limit-per-username—Maximum number of sessions allowed for a username per access profile.

Use the **brief** option to display only active users with blocked requests. Use the **detail** option to display all active users.

Required Privilege Level

view

RELATED DOCUMENTATION

Verifying and Managing Subscriber AAA Information

Understanding Session Options for Subscriber Access

Limiting the Number of Active Sessions per Username and Access Profile

clear network-access aaa statistics

List of Sample Output

[show network-access aaa statistics accounting on page 813](#)

[show network-access aaa statistics accounting detail on page 813](#)

[show network-access aaa statistics address-assignment client on page 814](#)

[show network-access aaa statistics address-assignment pool on page 815](#)

[show network-access aaa statistics address-assignment pool \(Excluded Addresses\) on page 815](#)

[show network-access aaa statistics dynamic-requests on page 815](#)

[show network-access aaa statistics radius on page 816](#)

[show network-access aaa statistics session-limit-per-username \(Users with Blocked Requests\) on page 816](#)

[show network-access aaa statistics session-limit-per-username \(All Active Users\) on page 816](#)

[show network-access aaa statistics session-limit-per-username on page 816](#)

Output Fields

[Table 36 on page 805](#) 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 36: show network-access aaa statistics Output Fields

Field Name	Field Description	Level of Output
Requests received	<ul style="list-style-type: none"> Number of accounting requests generated by the AAA framework. Number of dynamic requests received from the external server. <p>Does not include requests sent from backup accounting.</p>	All levels
Accounting request failures	<p>Number of accounting requests that failed to be sent or queued from a client to a RADIUS accounting server.</p> <p>Does not include requests sent from backup accounting.</p>	detail
Accounting request success	<p>Number of accounting requests successfully sent or queued from a client to a RADIUS accounting server.</p> <p>Does not include requests sent from backup accounting.</p>	detail

Table 36: show network-access aaa statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
Account on requests	Number of accounting on requests sent from a client to a RADIUS accounting server.	detail
Accounting start requests	Number of accounting start requests sent from a client to a RADIUS accounting server.	detail
Accounting interim requests	Number of accounting interim requests sent from a client to a RADIUS accounting server.	detail
Accounting stop requests	Number of accounting stop requests sent from a client to a RADIUS accounting server. Does not include requests sent from backup accounting.	detail
Accounting request timeouts	Number of accounting requests to the accounting server that timed out. This field was named Timed out requests in releases before Junos OS Release 16.1. Does not include requests sent from backup accounting.	All levels
Accounting Response failures	Number of accounting requests not acknowledged (NAK) by the accounting server. Does not include requests sent from backup accounting.	All levels
Accounting response success	Number of accounting requests acknowledged by the accounting server. Does not include requests sent from backup accounting.	All levels
Account on responses	Number of accounting on requests acknowledged by the RADIUS accounting server.	detail
Accounting start responses	Number of accounting start requests acknowledged by the RADIUS accounting server.	detail

Table 36: show network-access aaa statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
Accounting interim responses	Number of accounting interim requests acknowledged by the RADIUS accounting server.	detail
Accounting stop responses	Number of accounting stop requests acknowledged by the RADIUS accounting server. Does not include requests sent from backup accounting.	detail
Accounting rollover requests	Number of accounting requests coming to a RADIUS accounting server after a previous server timing out.	detail
Accounting unknown requests	Number of unknown accounting requests sent from a client to a RADIUS accounting server (for example, when the header has invalid or unsupported information).	detail
Accounting radius pending requests	Number of accounting requests sent from a client to a RADIUS accounting server that are waiting for a response from the server.	detail
Accounting malformed responses	Number of accounting responses from a RADIUS accounting server that have invalid or unexpected attributes.	detail
Accounting retransmissions	Number of accounting requests made by a client to the RADIUS sever that were retransmitted. Does not include requests sent from backup accounting.	detail
Accounting bad authenticators	Number of accounting responses from a RADIUS accounting server that have an incorrect authenticator (for example, the client and server RADIUS secret do not match).	detail
Accounting packets dropped	Number of accounting responses from a RADIUS accounting server that are dropped by a client.	detail

Table 36: show network-access aaa statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
Accounting backup record creation requests	Number of accounting stop requests from a client to a RADIUS accounting server that were forwarded to be backed up.	detail
Accounting backup replay request success	Number of backup accounting stop requests successfully created by clients after each timeout for replay to a RADIUS accounting server.	detail
Accounting backup request failures	Number of backup accounting requests that failed to be sent or queued from a client to a RADIUS accounting server.	detail
Accounting backup request success	Number of backup accounting requests successfully sent or queued from a client to a RADIUS accounting server.	detail
Accounting backup timeouts	Number of backup accounting requests that timed out after being sent to a RADIUS accounting server.	detail
Accounting backup in-flight requests	<p>Number of backup accounting requests that were successfully sent or queued to a RADIUS accounting server for which no response or error has been received yet.</p> <p>Backup requests are replayed only in the following circumstances:</p> <ul style="list-style-type: none"> • When the request being replayed receives a positive response, the next request can be replayed. • When the request being replayed receives a timeout response, it can be replayed again. <p>Consequently this intermediate timer displays 1 or 0. The value eventually drops to 0 as requests are responded to positively or fail due to error.</p>	detail

Table 36: show network-access aaa statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
Accounting backup responses success	Number of backup records that were successfully acknowledged with a positive response from a RADIUS accounting server.	detail
Accounting backup radius requests	Number of backup requests sent to UDP level. This is a RADIUS-level counter and increments rapidly based on the configured retries and timeouts and the RADIUS-level retransmissions. An observation that the value is increasing is more significant than the exact value of the counter.	detail
Accounting backup radius responses	Number of responses received at the UDP level for backup requests. This is a RADIUS-level counter and increments rapidly based on the configured retries and timeouts and the RADIUS-level retransmissions. Observation that the value is increasing is more significant than the exact value of the counter.	detail
Accounting backup radius timeouts	Number of backup requests that timed out after being sent to UDP. This is a RADIUS-level counter and increments rapidly based on the configured retries and timeouts and the RADIUS-level retransmissions. Observation that the value is increasing is more significant than the exact value of the counter.	detail
Accounting backup radius pending requests	Number of backup requests sent to a RADIUS accounting server that are waiting for a response from the server. This is an intermediate state counter that eventually drops to zero as requests are responded to or failed due to error.	detail

Table 36: show network-access aaa statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
Accounting backup radius retransmissions	Sum of backup request retransmissions for each RADIUS accounting server. This is a RADIUS-level counter and increments rapidly based on the configured retries and timeouts and the RADIUS-level retransmissions. Observation that the value is increasing is more significant than the exact value of the counter.	detail
Accounting backup malformed responses	Sum of malformed responses received for backup requests sent to each RADIUS accounting server at the UDP level.	detail
Accounting backup bad authenticators	Sum of responses received for backup accounting requests for each RADIUS accounting server where authenticators were mismatched.	detail
Accounting backup responses dropped	Sum of responses for backup accounting requests for each RADIUS accounting server that were dropped due to various sanity checks.	detail
Accounting backup rollover requests	Sum of backup accounting requests rolled over for each RADIUS accounting server.	detail
Accounting backup unknown responses	Sum of unknown responses for backup accounting requests for each RADIUS accounting server.	detail
Client	Client type; for example, DHCP, Mobile IP, PPP.	none specified
Out of Memory	Number of times an address was not given to the client due to memory issues.	none specified
No Matches	Number of times there were no network matches for the pool.	none specified

Table 36: show network-access aaa statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
Pool Name	Name of the address-assignment pool for this client.	none specified
Out of Addresses	Number of times there were no available addresses in the pool.	none specified
Address total	Number of addresses in the pool.	none specified
Addresses in use	Number of addresses in use.	none specified
Addresses excluded	Number of addresses excluded from being allocated from the pool with the excluded-address or excluded-range statements.	none specified
Address Usage (percent)	Percentage of total addresses in use. This value does not take excluded addresses into account.	none specified
Pool drain configured	Configuration state of active drain for the specified local address pool, yes or no .	none specified
Pool Usage	Percentage of allocated addresses in the specified address pool.	none specified
processed successfully	Number of dynamic requests processed successfully by the AAA framework.	All levels
errors during processing	Number of dynamic requests that resulted in processing errors by the AAA framework.	All levels
Link Name	Name of the secondary address-assignment pool to which the primary pool is linked.	
silently dropped	Number of dynamic requests dropped by the AAA framework due to multiple back-to-back or duplicate requests.	All levels
RADIUS Server	IPv4 or IPv6 address of the RADIUS server to which the router is sending requests.	All levels

Table 36: show network-access aaa statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
Profile	Name of the RADIUS profile associated with the RADIUS server. A RADIUS server can be associated with more than one RADIUS profile.	All levels
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.	All levels
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.	All levels
Peak	<p>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.</p> <p>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.</p>	All levels
Exceeded	<p>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.</p> <p>NOTE: If the value of this field is nonzero, you may want to increase the value of the Configured field.</p>	All levels
Username	Username for a subscriber with one or more active sessions for an access profile.	brief detail
Access-profile	Name of the access profile where the username is active.	brief detail

Table 36: show network-access aaa statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
Blocked requests	Number of session requests that have been blocked for the username for an access profile. A request is blocked when it exceeds the configured session limit.	brief detail
Session count	Number of active sessions for the username for an access profile.	brief detail
Total usernames	Number of active usernames for all access profiles.	none summary
Total usernames exceeding session limit	Number of usernames that have attempted sessions greater than the limit configured for the username.	none summary
Total blocked requests	Number of session requests that have been blocked because the session limit is exceeded.	none summary

Sample Output

show network-access aaa statistics accounting

user@host> **show network-access aaa statistics accounting**

```
Accounting module statistics
Accounting module statistics
  Requests received: 5000
  Accounting request timeouts: 2000
  Accounting response failures: 0
  Accounting response success: 3000
```

show network-access aaa statistics accounting detail

user@host> **show network-access aaa statistics accounting detail**

```

Accounting module statistics
Accounting module statistics
  Requests received: 5000
  Accounting request failures: 0
  Accounting request success: 5000
    Account on requests: 0
    Accounting start requests: 3000
    Accounting interim requests: 0
    Accounting stop requests: 2000
  Accounting request timeouts: 2000
  Accounting response failures: 0
  Accounting response success: 3000
    Account on responses: 0
    Accounting start responses: 3000
    Accounting interim responses: 0
    Accounting stop responses: 0
  Accounting rollover requests: 0
  Accounting unknown responses: 0
  Accounting radius pending requests: 0
  Accounting malformed responses: 0
  Accounting retransmissions: 6000
  Accounting bad authenticators: 0
  Accounting packets dropped: 0

  Accounting backup record creation requests: 3000
  Accounting backup request replay success: 9808
  Accounting backup request failures: 0
  Accounting backup request success: 3006
  Accounting backup timeouts: 6
  Accounting backup in-flight requests: 0
  Accounting backup responses success: 3000
  Accounting backup radius requests: 3006
  Accounting backup radius responses: 3000
  Accounting backup radius timeouts: 99
  Accounting backup radius pending requests: 0
  Accounting backup radius retransmissions: 99
  Accounting backup malformed responses: 0
  Accounting backup bad authenticators: 0
  Accounting backup responses dropped: 0
  Accounting backup rollover requests: 0
  Accounting backup unknown responses: 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
Pool drain configured: yes
```

show network-access aaa statistics address-assignment pool (Excluded Addresses)

```
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: 0
Address total: 24000
Addresses in use: 12000
Addresses excluded: 1000
Address Usage (percent): 50
Pool drain configured: yes
```

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
198.51.100.239	prof1	1000	0	1000	14
	prof2	500	17	432	0
198.51.100.211	myprof	200	0	200	27
203.0.113.254	pppoe-auth	111	0	1	0
2001:db8:0:f101::2	xyz-profile11	1000	10	135	0

show network-access aaa statistics session-limit-per-username (Users with Blocked Requests)

```
user@host> show network-access aaa statistics session-limit-per-username brief
```

Username	Access-profile	Blocked requests	Session count
xyz@example.net	BNG1	3	5
abc@example.net	BNG2	2	5

show network-access aaa statistics session-limit-per-username (All Active Users)

```
user@host> show network-access aaa statistics session-limit-per-username detail
```

Username	Access-profile	Blocked requests	Session count
rkv@example.net	BNG1	0	4
xyz@example.net	BNG1	3	5
abc@example.net	BNG2	2	5
pqr@example.net	BNG2	0	1

show network-access aaa statistics session-limit-per-username

```
user@host> show network-access aaa statistics on-limit-per-username
```

```
Total usernames: 15
Total usernames exceeding session limit: 2
Total blocked requests: 5
```

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 819](#)
[show network-access aaa statistics authentication detail on page 820](#)

Output Fields

Table 37 on page 817 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 37: 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
Accepts	Number of authentication requests accepted by the authentication server.	All levels

Table 37: show network-access aaa statistics authentication Output Fields (*continued*)

Field Name	Field Description	Level of Output
Rejects	Number of authentication requests rejected by the authentication server.	All levels
Challenges	Number of authentication requests challenged by the authentication server.	All levels
Timed out requests	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
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

Table 37: show network-access aaa statistics authentication Output Fields (*continued*)

Field Name	Field Description	Level of Output
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
Accepts: 261
Rejects: 975
```

```
Challenges: 0  
Timed out requests: 882
```

show network-access aaa statistics authentication detail

user@host> **show network-access aaa statistics authentication detail**

```
Authentication module statistics  
Requests received: 2118  
  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  
Timed out requests: 882
```


show network-access aaa subscribers

Syntax

```
show network-access aaa subscribers  
  <logical-system logical-system-name>  
  <routing-instance routing-instance-name>  
  <statistics>  
  <username>  
  <session-id session-id-number detail>
```

Release Information

Command introduced in Junos OS Release 9.1.

Command updated with **session-id *session-id-number* detail** in Junos OS Release 17.3.

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.

session-id *session-id-number* detail—(Optional) Display information for the specified session ID.

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 824](#)

[show network-access aaa subscribers logical-system routing-instance on page 824](#)

[show network-access aaa subscribers statistics username on page 824](#)

[show network-access aaa subscribers username on page 825](#)

[show network-access aaa subscribers session-id 26 detail on page 825](#)

Output Fields

Table 38 on page 822 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 38: 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.
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.

Table 38: show network-access aaa subscribers Output Fields (*continued*)

Field Name	Field Description
Session-ID	ID of the subscriber session.
Session uptime	How long the session has been up, in <i>HH:MM:SS</i> .
Accounting	Status of accounting, and type of accounting if accounting is on.
Stripped username	Username of the subscriber session.
AAA Logical system/Routing instance	AAA framework for the subscriber of logical system or routing instance.
Target Logical system/Routing instance	Target framework for the subscriber of logical system or routing instance.
Access-profile	Profile of the subscriber.
Accounting Session ID	ID of the subscriber session for accounting.
Multi Accounting Session ID	ID of the subscriber session for multiple accounting.
IP Address	IPv4 address of the subscriber.
IPv6 Address	IPv6 address of the subscriber.
IPv6 Prefix	IPv6 prefix of the subscriber.
Authentication State	State of subscriber session authentication.
Accounting State	State of subscriber session accounting.
Provisioning Type	Type of subscriber provisioning.

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
user61@example.net	ppp	default
00010e020304.1231	dhcp	isp-bos-metro-12:isp-cmbrg-12
user54@example.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
user54@example.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
  Successfull 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 user80@example.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 aaa subscribers session-id 26 detail

The following command output is seen when only an IPv4 client is associated with the session:

user@host> **show network-access aaa subscribers session-id 26 detail**

```

Type: dhcp
Stripped username: my-customer
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:default
Access-profile: AccessProfile
Session ID: 26
Accounting Session ID: 26
Multi Accounting Session ID: 0
IP Address: 10.0.0.2
Authentication State: AuthStateActive
Accounting State: Acc-Interim-Sent
Provisioning Type: None

```

The following command output is seen when IPv6 client logs in (after IPv4 association) and is associated with the same session ID:

user@host> **show network-access aaa subscribers session-id 26 detail**

```
Type: dhcp
Stripped username: my-customer
AAA Logical system/Routing instance: default:default
Target Logical system/Routing instance: default:default
Access-profile: AccessProfile
Session ID: 26
Accounting Session ID: 26
Multi Accounting Session ID: 0
IP Address: 10.0.0.2
IPv6 Address: 2001:db8:0:8003::2
IPv6 Prefix: 2001:db8:ffff:0:4::/64
Authentication State: AuthStateActive
Accounting State: Acc-Interim-Sent
Provisioning Type: None
```

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 828](#)

Output Fields

[Table 39 on page 827](#) lists the output fields for the **show network-access address-assignment pool** command. Output fields are listed in the approximate order in which they appear.

Table 39: show network-access address-assignment pool Output Fields

Field Name	Field Description
IP address/prefix	IP address of the client.
Hardware address	MAC address of the client. Displays NA for addresses excluded from being allocated from the pool with the excluded-address or excluded-range statements.

Table 39: show network-access address-assignment pool Output Fields (*continued*)

Field Name	Field Description
Host/User	<p>Hostname or username of the client.</p> <p>Displays EXCLUDED for addresses excluded from being allocated from the pool with the excluded-address or excluded-range statements.</p>
Type	<p>Type of client.</p> <p>Displays unknown for addresses excluded from being allocated from the pool with the excluded-address or excluded-range statements.</p>

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/prefix	Hardware address	Host/User	Type
192.168.2.1	00:00:5e:00:53:01	user1	DHCP
192.168.2.2	00:00:5e:00:53:02	user2	DHCP
192.168.2.3	00:00:5e:00:53:03	user3	DHCP
192.168.2.4	NA	EXCLUDED	unknown

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.

Starting in Junos OS Release 17.3, the * (asterisk) wildcard character is supported for the interface name for debugging purpose. With this support, you can match any string of characters in that position in the interface name. For example, so* matches all SONET/SDH interfaces.

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

Required Privilege Level

view

List of Sample Output

- [show ppp interface on page 846](#)
- [show ppp interface extensive \(LCP Connection Update Negotiation Successful\) on page 846](#)
- [show ppp interface extensive \(LCP Connection Update Negotiation Failed\) on page 847](#)
- [show ppp interface extensive \(Inline Service Interface\) on page 847](#)
- [show ppp interface terse on page 848](#)

Output Fields

[Table 40 on page 829](#) lists the output fields for the **show ppp interface** command. Output fields are listed in the approximate order in which they appear.

Table 40: 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

Table 40: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
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
<i>protocol State</i>	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
Keepalive settings	<p>Keepalive settings for the PPP sessions on the L2TP network server (LNS). LNS-based PPP sessions are supported only on service interfaces (si).</p> <ul style="list-style-type: none"> • 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
Magic-Number validation	<p>Indicates whether the local peer is configured to ignore mismatches between peer magic numbers when the numbers are validated during PPP keepalive (Echo-Request/Echo-Reply) exchanges.</p> <ul style="list-style-type: none"> • Enable—Mismatch detection sends failed Echo-Reply packets to the Routing Engine. If a valid magic number is not received within the configurable keepalive interval, PPP treats this as a keepalive failure and tears down the PPP sessions. • Disable—The Packet Forwarding Engine does not perform a validation check for magic numbers received from remote peers. A mismatch cannot be detected, so receipt of its own magic number or an unexpected value does not trigger notification to the Routing Engine. 	extensive

Table 40: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
RE Keepalive statistics	<p>Keepalive statistics for the packets handled by the Routing Engine.</p> <ul style="list-style-type: none"> • 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 40: show ppp interface Output Fields (continued)

Field Name	Field Description	Level of Output
LCP		extensive

Table 40: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
	<p>LCP information:</p> <ul style="list-style-type: none"> • State—LCP protocol state (all platforms except M120 and M320 routers): <ul style="list-style-type: none"> • 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): <ul style="list-style-type: none"> • 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. • Last updated—Reports the timestamp of the last successful connection update exchange. <ol style="list-style-type: none"> 1. When LCP negotiation completes, this field has the same value as the Last completed field. 2. The field then reports the timestamp of any subsequent successful exchange of Connection-Update-Request and Connection-Update-Ack messages with 	

Table 40: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
	<p>the peer (such as a home gateway).</p> <p>This field is displayed only when the Connection-Status-Message option is successfully negotiated.</p>	

Table 40: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
------------	-------------------	-----------------

Table 40: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
	<ul style="list-style-type: none"> • Negotiated options: <ul style="list-style-type: none"> • 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. • Connection Update Requests—Number of connection update requests sent by PPP to the remote peer (such as a home gateway). This value does not include retries. This field is displayed even when negotiation fails for the Connection-Status-Message option. This enables you to confirm that an update request was sent. The absence of the Juniper Connection Status Message field indicates the peer does not support the updates. • 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. • Juniper Connection Status Message—The content of the message conveyed by RADIUS in the Connection-Status-Message VSA (26-4874-218). This field is displayed only when the Connection-Status-Message option is successfully negotiated. • 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. 	

Table 40: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
	<ul style="list-style-type: none"> • 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 40: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Authentication	<p>CHAP or PAP authentication state information. For CHAP authentication:</p> <ul style="list-style-type: none"> • 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. <p>For PAP authentication:</p> <ul style="list-style-type: none"> • Pap-resp-sent—PAP response sent to peer (ACK/NACK). • 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 40: show ppp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
IPCP		extensive

Table 40: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
	<p>Internet Protocol Control Protocol (IPCP) information.</p> <ul style="list-style-type: none"> • State—(All platforms except M120 and M320 routers) One of the following values: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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. 	

Table 40: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
	<ul style="list-style-type: none"> • 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. • Negotiation mode—PPP Network Control Protocol (NCP) negotiation mode configured for IPCP: Active or Passive 	

Table 40: show ppp interface Output Fields (continued)

Field Name	Field Description	Level of Output
IPV6CP		extensive

Table 40: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
	<p>Internet Protocol version 6 Control Protocol (IPv6CP) information.</p> <ul style="list-style-type: none"> • State—(All platforms except M120 and M320 routers) One of the following values: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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 	

Table 40: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
	<p>quad notation.</p> <ul style="list-style-type: none"> • Negotiation mode—PPP Network Control Protocol (NCP) negotiation mode configured for IPv6CP: Active or Passive 	
OSINLCP State	<p>OSI Network Layer Control Protocol (OSINLCP) protocol state information (all platforms except M120 and M320 routers):</p> <ul style="list-style-type: none"> • State: <ul style="list-style-type: none"> • 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—OSINLCP state completion time. 	extensive

Table 40: show ppp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
TAGCP	<p>TAGCP information.</p> <ul style="list-style-type: none"> • State—(All platforms except M120 and M320 routers) One of the following values: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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 (LCP Connection Update Negotiation Successful)

```
user@host> show ppp interface extensive pp0.3221225489
```

```
Session pp0.3221225489, Type: PPP, Phase: Network
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
    Magic-Number validation: enable
LCP
    State: Opened
    Last started: 2020-02-11 15:06:00 PDT
    Last completed: 2020-02-11 15:06:00 PDT
    Last updated: 2020-02-11 15:06:10 PDT
    Negotiated options:
        Magic number: 906403799, Initial Advertised MRU: 1492, Local MRU: 1492,
Peer MRU: 149
        Juniper Connection Status Message: 10m:xxxx
        Connection-Update-Requests: 1
    Authentication: Off
IPCP
    State: Opened
    Last started: 2020-02-11 15:06:00 PDT
    Last completed: 2020-02-11 15:06:00 PDT
    Negotiated options:
        Local address: 198.51.100.30, Remote address: 203.0.113.9
    Negotiation mode: Passive
IPV6CP
    State: Opened
    Last started: 2020-02-11 15:06:00 PDT
    Last completed: 2020-02-11 15:06:00 PDT
    Negotiated options:
        Local interface identifier: 2001:db8::fc73:cba, Remote interface identifier:
```

```
2001:db8::3a
Negotiation mode: Passive
```

show ppp interface extensive (LCP Connection Update Negotiation Failed)

```
user@host> show ppp interface extensive pp0.3221225489
```

```
Session pp0.3221225489, Type: PPP, Phase: Network
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
    Magic-Number validation: enable

LCP
  State: Opened
  Last started: 2020-02-11 15:06:00 PDT
  Last completed: 2020-02-11 15:06:00 PDT
  Negotiated options:
    Magic number: 906403799, Initial Advertised MRU: 1492, Local MRU: 1492,
Peer MRU: 149
    Connection-Update-Requests: 1
  Authentication: Off

IPCP
  State: Opened
  Last started: 2020-02-11 15:06:00 PDT
  Last completed: 2020-02-11 15:06:00 PDT
  Negotiated options:
    Local address: 198.51.100.30, Remote address: 203.0.113.9
  Negotiation mode: Passive

IPV6CP
  State: Opened
  Last started: 2020-02-11 15:06:00 PDT
  Last completed: 2020-02-11 15:06:00 PDT
  Negotiated options:
    Local interface identifier: 2001:db8::fc73:cba, Remote interface identifier:
2001:db8::3a
  Negotiation mode: Passive
```

show ppp interface extensive (Inline Service Interface)

```
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
    Magic-Number validation: disable

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 timeout : 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: 203.0.113.21, Remote address: 203.0.113.22
  Negotiation mode: Active
IPV6CP
  State: Opened
  Last started: 2007-01-29 10:43:50 PST
  Last completed: 2007-01-29 10:43:50 PST
  Negotiated options:
    Local interface identifier: 2a0:a522:64:d319, Remote interface identifier:
0:0:0:c
  Negotiation mode: Passive

```

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

show subscribers

Syntax

```
show subscribers
<detail | extensive | terse>
<aci-interface-set-name aci-interface-set-name>
<address address>
<agent-circuit-identifier agent-circuit-identifier>
<agent-remote-identifier agent-remote-identifier>
<aggregation-interface-set-name interface-set-name>
<client-type client-type>
<count>
<id session-id <accounting-statistics>>
<interface interface <accounting-statistics>>
<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 12.3Rx releases.

Options **vci** and **vpi** supported in Junos OS Release 13.2 and later releases. (Not supported in Junos OS Release 13.1.)

Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Enhanced subscriber management supported in Junos OS Release 15.1R3 on MX Series routers.

accounting-statistics option added in Junos OS Release 15.1R3 and 17.4R1 on MX Series routers.

aggregation-interface-set-name option added in Junos OS Release 18.4R1 on MX Series routers.

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, 192.0.2.0). If you specify the IP address as a prefix with a netmask (for example, 192.0.2.0/32), the router displays a message that the IP address is invalid, and rejects the command.

agent-circuit-identifier—(Optional) Display all dynamic subscriber sessions whose ACI value matches the specified string. You can specify either the complete ACI string or a substring. To specify a substring, you must enter characters that form the beginning of the string, followed by an asterisk (*) as a wildcard to substitute for the remainder of the string. The wildcard can be used only at the end of the specified substring; for example:

```
user@host1> show subscribers agent-circuit-identifier substring*
```

Junos OS Release	Substring Support
Junos OS Release 13.3R1	You can specify a substring without a wildcard.
Starting in Junos OS Release 14.1R1	You must specify the complete ACI string; you cannot specify a wildcard.
Starting in Junos OS Release 15.1R7, 16.1R7, 16.2R3, 17.1R3, 17.2R3, 17.3R3, 17.4R2, 18.1R2, 18.2R1	You can specify a substring, but you must include the wildcard character at the end of the substring.

agent-remote-identifier—(Optional) Display all dynamic subscriber sessions whose ARI value matches the specified string. You must specify the complete ACI string; you cannot specify a wildcard.

aggregation-interface-set-name interface-set-name—(Optional) Display summary information for the specified aggregation node interface set, including interface, VLAN ID, username and LS:RI.

client-type—(Optional) Display subscribers whose client type matches one of the following client types:

- **dhcp**—DHCP clients only.
- **dotlx**—Dotlx clients only.
- **essm**—ESSM clients only.

- **fixed-wireless-access**—Fixed wireless access clients only.
- **fwauth**—FWAuth (authenticated across a firewall) clients only.
- **l2tp**—L2TP clients only.
- **mlppp**—MLPPP clients only.
- **ppp**—PPP clients only.
- **pppoe**—PPPoE clients only.
- **static**—Static clients only.
- **vlan**—VLAN clients only.
- **vlan-oob**—VLAN out-of-band (ANCP-triggered) clients only.
- **vpls-pw**—VPLS pseudowire clients only.
- **xauth**—Xauth clients only.

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 session-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.

id session-id accounting-statistics—(Optional) Display accurate subscriber accounting statistics for a subscriber session with the specified ID. Requires the **actual-transmit-statistics** statement to be configured in the dynamic profile for the dynamic logical interface. If the statement is not configured, a value of 0 is displayed for accounting statistics.

interface—(Optional) Display subscribers whose interface matches the specified interface.

interface accounting-statistics—(Optional) Display subscriber accounting statistics for the specified interface. Requires the **actual-transmit-statistics** statement to be configured in the dynamic profile for the dynamic logical 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.

stacked-vlan-id—(Optional) Display subscribers whose stacked VLAN ID matches the specified stacked VLAN ID.

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 65,535.

vlan-id—(Optional) Display subscribers whose VLAN ID matches the specified VLAN ID, regardless of whether the subscriber uses a single-tagged or double-tagged VLAN. For subscribers using a double-tagged VLAN, this option displays subscribers where the inner VLAN tag matches the specified VLAN ID. To display only subscribers where the specified value matches only double-tagged VLANs, use the **stacked-vlan-id** **stacked-vlan-id** option to match the outer VLAN tag.

NOTE: Because of display limitations, logical system and routing instance output values are truncated when necessary.

Required Privilege Level

view

RELATED DOCUMENTATION

[show subscribers summary](#) | 896

Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration

Verifying and Managing Configurations for Dynamic VLANs Based on Access-Line Identifiers

Verifying and Managing Junos OS Enhanced Subscriber Management

List of Sample Output

[show subscribers \(IPv4\) on page 864](#)

[show subscribers \(IPv6\) on page 864](#)

[show subscribers \(IPv4 and IPv6 Dual Stack\) on page 864](#)

[show subscribers \(Single Session DHCP Dual Stack\) on page 864](#)

[show subscribers \(Single Session DHCP Dual Stack detail\) on page 865](#)
[show subscribers \(LNS on MX Series Routers\) on page 865](#)
[show subscribers \(L2TP Switched Tunnels\) on page 866](#)
[show subscribers aggregation-interface-set-name on page 866](#)
[show subscribers client-type dhcp detail on page 866](#)
[show subscribers client-type dhcp detail \(DHCPv6\) on page 867](#)
[show subscribers client-type dhcp extensive on page 868](#)
[show subscribers client-type fixed-wireless-access on page 868](#)
[show subscribers client-type fixed-wireless-access detail \(Detail\) on page 869](#)
[show subscribers client-type vlan-oob detail on page 869](#)
[show subscribers count on page 870](#)
[show subscribers address detail \(IPv6\) on page 870](#)
[show subscribers detail \(IPv4\) on page 871](#)
[show subscribers detail \(IPv6\) on page 871](#)
[show subscribers detail \(pseudowire Interface for GRE Tunnel\) on page 872](#)
[show subscribers detail \(IPv6 Static Demux Interface\) on page 872](#)
[show subscribers detail \(L2TP LNS Subscribers on MX Series Routers\) on page 873](#)
[show subscribers detail \(L2TP Switched Tunnels\) on page 873](#)
[show subscribers detail \(Tunneled Subscriber\) on page 874](#)
[show subscribers detail \(IPv4 and IPv6 Dual Stack\) on page 874](#)
[show subscribers detail \(ACI Interface Set Session\) on page 875](#)
[show subscribers detail \(PPPoE Subscriber Session with ACI Interface Set\) on page 876](#)
[show subscribers extensive on page 876](#)
[show subscribers extensive \(Aggregation Node Interface Set and DSL Forum Attributes\) on page 877](#)
[show subscribers extensive \(Passive Optical Network Circuit Interface Set\) on page 878](#)
[show subscribers extensive \(DNS Addresses from Access Profile or Global Configuration\) on page 879](#)
[show subscribers extensive \(DNS Addresses from RADIUS\) on page 880](#)
[show subscribers extensive \(IPv4 DNS Addresses from RADIUS, IPv6 from Access Profile or Global Configuration\) on page 880](#)
[show subscribers extensive \(RPF Check Fail Filter\) on page 881](#)
[show subscribers extensive \(L2TP LNS Subscribers on MX Series Routers\) on page 881](#)
[show subscribers extensive \(IPv4 and IPv6 Dual Stack\) on page 882](#)
[show subscribers extensive \(ADF Rules \) on page 883](#)
[show subscribers extensive \(Effective Shaping-Rate\) on page 884](#)
[show subscribers extensive \(PPPoE Subscriber Access Line Rates on page 884](#)
[show subscribers extensive \(Subscriber Session Using PCEF Profile\) on page 886](#)
[show subscribers aci-interface-set-name detail \(Subscriber Sessions Using Specified ACI Interface Set\) on page 887](#)
[show subscribers agent-circuit-identifier detail \(Subscriber Sessions Using Specified ACI Substring\) on page 888](#)
[show subscribers id accounting-statistics on page 889](#)
[show subscribers interface accounting-statistics on page 889](#)
[show subscribers interface extensive on page 890](#)

[show subscribers logical-system terse on page 891](#)
[show subscribers physical-interface count on page 891](#)
[show subscribers routing-instance inst1 count on page 892](#)
[show subscribers stacked-vlan-id detail on page 892](#)
[show subscribers stacked-vlan-id vlan-id detail \(Combined Output\) on page 892](#)
[show subscribers stacked-vlan-id vlan-id interface detail \(Combined Output for a Specific Interface\) on page 892](#)
[show subscribers user-name detail on page 893](#)
[show subscribers vlan-id on page 893](#)
[show subscribers vlan-id detail on page 893](#)
[show subscribers vpi vci extensive \(PPPoE-over-ATM Subscriber Session\) on page 894](#)
[show subscribers address detail \(Enhanced Subscriber Management\) on page 894](#)

Output Fields

Table 41 on page 854 lists the output fields for the **show subscribers** command. Output fields are listed in the approximate order in which they appear.

Table 41: 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 <i>tpid.vlan-id</i> 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, FWA, GRE, L2TP, PPP, PPPoE, STATIC-INTERFACE, VLAN).
IP Address	Subscriber IPv4 address.

Table 41: show subscribers Output Fields (*continued*)

Field Name	Field Description
IP Netmask	<p>Subscriber IP netmask.</p> <p>(MX Series) This field displays 255.255.255.255 by default. For tunneled or terminated PPP subscribers only, this field displays the actual value of Framed-IP-Netmask when the SDB_FRAMED_PROTOCOL attribute in the session database is equal to AUTHD_FRAMED_PROTOCOL_PPP. This occurs in the use case where the LNS generates access-internal routes when it receives Framed-IP-Netmask from RADIUS during authorization. When it receives Framed-Pool from RADIUS, the pool mask is ignored and the default /32 mask is used.</p>
Primary DNS Address	<p>IP address of primary DNS server.</p> <p>This field is displayed with the extensive option only when the address is provided by RADIUS.</p>
Secondary DNS Address	<p>IP address of secondary DNS server.</p> <p>This field is displayed with the extensive option only when the address is provided by RADIUS.</p>
IPv6 Primary DNS Address	<p>IPv6 address of primary DNS server.</p> <p>This field is displayed with the extensive option only when the address is provided by RADIUS.</p>
IPv6 Secondary DNS Address	<p>IPv6 address of secondary DNS server.</p> <p>This field is displayed with the extensive option only when the address is provided by RADIUS.</p>
Domain name server inet	<p>IP addresses for the DNS server, displayed in order of configuration.</p> <p>This field is displayed with the extensive option only when the addresses are derived from the access profile or the global access configuration.</p>
Domain name server inet6	<p>IPv6 addresses for the DNS server, displayed in order of configuration.</p> <p>This field is displayed with the extensive option only when the addresses are derived from the access profile or the global access configuration.</p>
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.

Table 41: show subscribers Output Fields (*continued*)

Field Name	Field Description
IPv6 Prefix	Subscriber IPv6 prefix. If you are using DHCPv6 prefix delegation, this is the delegated prefix.
IPv6 User Prefix	IPv6 prefix obtained through NDRA.
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.
Logical System	Logical system associated with the subscriber.
Routing Instance	Routing instance associated with the subscriber.
Interface	(Enhanced subscriber management for MX Series routers) Name of the enhanced subscriber management logical interface, in the form demux0.nnnn (for example, demux0.3221225472), to which access-internal and framed subscriber routes are mapped.
Interface Type	Whether the subscriber interface is Static or Dynamic .

Table 41: show subscribers Output Fields (*continued*)

Field Name	Field Description
Interface Set	<p>Internally generated name of the dynamic ACI or ALI interface set used by the subscriber session. The prefix of the name indicates the string received in DHCP or PPPoE control packets on which the interface set is based. For ALI interface sets, the prefix indicates that the value is configured as a trusted option to identify the subscriber line.</p> <p>The name of the interface set uses one of the following prefixes:</p> <ul style="list-style-type: none"> • aci—ACI; for example, aci-1033-demux0.3221225524. This is the only prefix allowed for ACI interface sets. • ari—ARI; for example, ari-1033-demux0.3221225524. • aci+ari—Both the ACI and ARI; for example, aci+ari-1033-demux0.3221225524. • noids—Neither the ACI nor the ARI were received; for example, noids-1033-demux0.3221225524. <p>NOTE: ACI interface sets are configured with the agent-circuit-identifier autoconfiguration stanza. ALI interface sets are configured with the line-identity autoconfiguration stanza.</p> <p>Besides dynamic ACI and ALI interface sets, this field can be an interface set based on a substring of the ARI string. This occurs when the dynamic profile includes the predefined variable \$junos-pon-id-interface-set-name, and the profile is applied for a passive optical network (PON). The ARI string is inserted by the optical line terminal (OLT). The final substring in the string, unique for the PON, identifies individual subscriber circuits, and is used as the name of the interface set.</p>
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).

Table 41: show subscribers Output Fields (*continued*)

Field Name	Field Description
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).
PFE Flow ID	Forwarding flow identifier.
VLAN Id	VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i> .
Stacked VLAN Id	Stacked VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i> .
RADIUS Accounting ID	RADIUS accounting ID associated with the subscriber.
Agent Circuit ID	<p>For the dhcp client type, 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.</p> <p>For the vlan-oob client type, the agent circuit ID or access-loop circuit identifier that identifies the subscriber line based on the subscriber-facing DSLAM interface on which the subscriber request originates.</p>
Agent Remote ID	<p>For the dhcp client type, 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.</p> <p>For the vlan-oob client type, the agent remote ID or access-loop remote identifier that identifies the subscriber line based on the NAS-facing DSLAM interface on which the subscriber request originates.</p>

Table 41: show subscribers Output Fields (continued)

Field Name	Field Description
Aggregation Interface-set Name	<p>Value of the \$junos-aggregation-interface-set-name predefined variable; one of the following:</p> <ul style="list-style-type: none"> When the hierarchical-access-network-detection option is configured for the access lines and the value of the Access-Aggregation-Circuit-ID-ASCII attribute (TLV 0x0003) received either in the ANCP Port Up message or PPPoE PADR IA tags begins with a # character, then the variable takes the value of the remainder of the string after the # character. When the hierarchical-access-network-detection option is not configured, or if the sting does not begin with the # character, then the variable takes the value specified with the predefined-variable-defaults statement.
Accounting Statistics	Actual transmitted subscriber accounting statistics by session ID or interface. Service accounting statistics are not included. These statistics do not include overhead bytes or dropped packets; they are the accurate statistics used by RADIUS. The statistics are counted when the actual-transmit-statistics statement is included in the dynamic profile.
DHCP Relay IP Address	IP address used by the DHCP relay agent.
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.
DHCPV6 Options	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCPv6 options.
Server DHCP Options	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCP options.
Server DHCPV6 Options	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCPv6 options.
DHCPV6 Header	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCPv6 options.
Effective shaping-rate	Actual downstream traffic shaping rate for the subscriber, in kilobits per second.

Table 41: show subscribers Output Fields (*continued*)

Field Name	Field Description
IPv4 Input Service Set	Input service set in access dynamic profile.
IPv4 Output Service Set	Output service set in access dynamic profile.
PCEF Profile	PCEF profile in access dynamic profile.
PCEF Rule/Rulebase	PCC rule or rulebase used in dynamic profile.
Dynamic configuration	Values for variables that are passed into the dynamic profile from RADIUS.
Service activation time	Time at which the first family in this service became active.
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 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 ID	ID number for a subscriber service session.
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.

Table 41: show subscribers Output Fields (*continued*)

Field Name	Field Description
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. This field is displayed only when the predefined variable \$junos-ipv6-address is used in the dynamic profile.
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.
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).

Table 41: show subscribers Output Fields (*continued*)

Field Name	Field Description
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).
DSL type	PPPoE subscriber's access line type reported by the PPPoE intermediate agent in a PADI or PADO packet in the Vendor-Specific-Tags TLV in subattribute DSL-Type (0x0091). The DSL type is one of the following types: ADSL , ADSL2 , ADSL2+ , OTHER , SDSL , VDSL , or VDSL2 .
Frame/Cell Mode	<p>Mode type of the PPPoE subscriber's access line determined by the PPPoE daemon based on the received subattribute DSL-Type (0x0091):</p> <ul style="list-style-type: none"> • Cell—When the DSL line type is one of the following: ADSL, ADSL2, or ADSL2+. • Frame—When the DSL line type is one of the following: OTHER, SDSL, VDSL, or VDSL2. <p>The value is stored in the subscriber session database.</p>
Overhead accounting bytes	Number of bytes added to or subtracted from the actual downstream cell or frame overhead to account for the technology overhead of the DSL line type. The value is determined by the PPPoE daemon based on the received subattribute DSL-Type (0x0091). The value is stored in the subscriber session database.
Actual upstream data rate	Unadjusted upstream data rate for the PPPoE subscriber's access line reported by the PPPoE intermediate agent in a PADI or PADO packet in the Vendor-Specific-Tags TLV in subattribute Actual-Net-Data-Rate-Upstream (0x0081).
Actual downstream data rate	Unadjusted downstream data rate for the PPPoE subscriber's access line reported by the PPPoE intermediate agent in a PADI or PADO packet in the Vendor-Specific-Tags TLV in subattribute Actual-Net-Data-Rate-Downstream (0x0082).
Adjusted downstream data rate	Adjusted downstream data rate for the PPPoE subscriber's access line, calculated by the PPPoE daemon and stored in the subscriber session database.
Adjusted upstream data rate	Adjusted upstream data rate for the PPPoE subscriber's access line, calculated by the PPPoE daemon and stored in the subscriber session database.

Table 41: show subscribers Output Fields (*continued*)

Field Name	Field Description
Local TEID-U	<p>Tunnel endpoint identifier on the BNG for the GTP-U user plane tunnel to the eNodeB. The identifier is allocated by the BNG.</p> <p>A fully qualified local TEID-C consists of this identifier and the GTPU Tunnel Local IP address value.</p>
Local TEID-C	<p>Tunnel endpoint identifier on the BNG for the GTP-C control plane tunnel to the MME. The identifier is allocated by the BNG.</p> <p>A fully qualified local TEID-C consists of this identifier and the GTPC Local IP address value.</p>
Remote TEID-U	<p>Tunnel endpoint identifier on the eNodeB for the GTP-U user plane tunnel to the BNG. The identifier is allocated by the eNodeB.</p> <p>A fully qualified remote TEID-U consists of this identifier and the GTPU Tunnel Remote IP address value.</p>
Remote TEID-C	<p>Tunnel endpoint identifier on the MME for the GTP-C control plane tunnel to the BNG. The identifier is allocated by the MME.</p> <p>A fully qualified remote TEID-C consists of this identifier and the GTPC Remote IP address value.</p>
GTPU Tunnel Remote IP address	<p>IP address of the S1-U interface on the eNodeB for the GTP-U tunnel endpoint.</p> <p>A fully qualified remote TEID-U consists of this address and the Remote TEID-U value.</p>
GTPU Tunnel Local IP address	<p>IP address of the S1-U interface on the BNG for the GTP-U tunnel endpoint.</p> <p>A fully qualified local TEID-U consists of this address and the Local TEID-U value</p>
GTPC Remote IP address	<p>IP address of the S11 interface on the MME for the GTP-C tunnel endpoint.</p> <p>A fully qualified remote TEID-C consists of this address and the Remote TEID-C value.</p>
GTPC Local IP address	<p>IP address of the S11 interface on the BNG for the GTP-C tunnel endpoint.</p> <p>A fully qualified local TEID-C consists of this address and the Local TEID-C value.</p>
Access Point Name	<p>Access point name (APN) for the user equipment. The APN corresponds to the connection and service parameters that the subscriber's mobile device can use for connecting to the carrier's gateway to the Internet.</p>

Sample Output

show subscribers (IPv4)

```
user@host> show subscribers
```

Interface	IP Address/VLAN ID	User Name	LS:RI
ge-1/3/0.1073741824	10		default:default
demux0.1073741824	203.0.113.10	WHOLESALE-CLIENT	default:default
demux0.1073741825	203.0.113.3	RETAILER1-CLIENT	test1:retailer1
demux0.1073741826	203.0.113.3	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:db8:c0:0:0:0/74	WHOLESALE-CLIENT	default:default
*	2001:db8:1/128	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	203.0.113.13	dualstackuser1@example1.com	
default:ASP-1			
*	2001:db8:1::/48		
*	2001:db8:1:1::/64		
pp0.1073741837	203.0.113.33	dualstackuser2@example1.com	
default:ASP-1			
*	2001:db8:1:2:5::/64		

show subscribers (Single Session DHCP Dual Stack)

```
user@host> show subscribers
```

Interface	IP Address/VLAN ID	User Name	LS:RI
demux0.1073741364	192.168.10.10	dual-stack-retail35	
default:default	2001:db8::100:0:0:0/74		
default:default	2001:db8:3ffe:0:4::/64		

show subscribers (Single Session DHCP Dual Stack detail)

user@host> show subscribers id 27 detail

```
Type: DHCP
User Name: dual-stack-retail33
IP Address: 10.10.0.53
IPv6 Address: 2001:db8:3000:0:0:8003::2
IPv6 Prefix: 2001:db8:3ffe:0:4::/64
Logical System: default
Routing Instance: default
Interface: ae0.3221225472
Interface type: Static
Underlying Interface: ae0.3221225472
Dynamic Profile Name: dhcp-retail-18
MAC Address: 00:00:5E:00:53:02
State: Active
DHCP Relay IP Address: 10.10.0.1
Radius Accounting ID: 27
Session ID: 27
PFE Flow ID: 2
Stacked VLAN Id: 2000
VLAN Id: 1
Login Time: 2014-05-15 10:12:10 PDT
DHCP Options: len 60
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 00 64 01 01 02
00 06 00 04 00 03 00 19 00 03 00 0c 00 00 00 00 00 00 00 00
00 00 00 00 00 19 00 0c 00 00 00 00 00 00 00 00 00 00 00 00
```

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	192.0.2.0	user@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	user@example.com	default:default
si-2/1/0.1073741843	Tunnel-switched	user@example.com	default:default

show subscribers aggregation-interface-set-name

```
user@host> show subscribers aggregation-interface-set-name FRA*
```

Interface	IP Address/VLAN ID	User Name
LS:RI		
ge-1/0/0.3221225472	50	ancp
default:ispl-subscriber		

show subscribers client-type dhcp detail

```
user@host> show subscribers client-type dhcp detail
```

```
Type: DHCP
IP Address: 203.0.113.29
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux
MAC Address: 00:00:5e:00:53:98
State: Active
Radius Accounting ID: user :2304
Login Time: 2009-08-25 14:43:52 PDT
```

```
Type: DHCP
IP Address: 203.0.113.27
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: 00:00:5e:00:53:f3
State: Active
Radius Accounting ID: 1234 :2560
Login Time: 2009-08-25 14:43:56 PDT

```

show subscribers client-type dhcp detail (DHCPv6)

user@host> show subscribers client-type dhcp detail

```

Type: DHCP
User Name: DEFAULTUSER
IPv6 Address: 2001:db8::2
IPv6 Prefix: 2001:db8:1::/64
Logical System: default
Routing Instance: default
Interface: demux0.3221225602
Interface type: Static
Underlying Interface: demux0.3221225602
Dynamic Profile Name: client-profile
MAC Address: 00:00:5E:00:53:01
State: Active
Radius Accounting ID: 142
Session ID: 142
PFE Flow ID: 148
Stacked VLAN Id: 1
VLAN Id: 1
Login Time: 2018-03-29 12:27:38 EDT
DHCP Options: len 56
00 08 00 02 00 00 00 01 00 0e 00 01 00 01 22 4f d0 33 00 11
01 00 00 01 00 03 00 0c 00 00 00 0a 00 04 9d 40 00 07 62 00
00 19 00 0c 00 00 00 0b 00 04 9d 40 00 07 62 00
Server DHCPV6 Options: len 94
00 0a 00 06 11 22 33 44 55 66 00 11 00 09 00 00 0c 4c 00 02
00 01 aa 00 11 00 20 00 00 0a 4c 00 02 00 02 32 33 00 03 00
03 34 35 36 00 05 00 06 31 32 33 34 35 36 00 06 00 01 31 00
11 00 09 00 00 0b 4c 00 02 00 01 bb 00 11 00 12 00 00 0d e9
00 01 00 03 aa bb cc 00 02 00 03 dd ee cc
DHCPV6 Header: len 4
01 fc e4 96

```

show subscribers client-type dhcp extensive

```
user@host> show subscribers client-type dhcp extensive
```

```
Type: DHCP
User Name: user
IP Address: 192.0.2.4
IP Netmask: 255.0.0.0
IPv6 Address: 2001:db8:3::103
IPv6 Prefix: 2001:db8::/68
Domain name server inet6: 2001:db8:1 abcd::2
Logical System: default
Routing Instance: default
Interface: ge-0/0/0.0
Interface type: Static
Underlying Interface: ge-0/0/0.0
MAC Address: 00:00:5e:00:53:01
State: Configured
Radius Accounting ID: 10
Session ID: 10
PFE Flow ID: 2
VLAN Id: 100
Agent Circuit ID: ge-0/0/0:100
Agent Remote ID: ge-0/0/0:100
Login Time: 2017-05-23 12:52:22 IST
DHCPV6 Options: len 69
00 01 00 0e 00 01 00 01 59 23 e3 31 00 10 94 00 00 01 00 08
00 02 00 00 00 19 00 29 00 00 00 00 00 04 9d 40 00 07 62 00
00 1a 00 19 00 09 3a 80 00 27 8d 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00
Server DHCP Options: len 13
3a 04 00 00 00 ff 00 3b 04 00 00 0f 00
Server DHCPV6 Options: len 8
00 0a 00 04 ab cd ef ab
DHCPV6 Header: len 4
01 00 00 04
IP Address Pool: al_pool30
IPv6 Address Pool: ia_na_pool
IPv6 Delegated Address Pool: prefix_delegate_pool
```

show subscribers client-type fixed-wireless-access

```
user@host> show subscribers client-type fixed-wireless-access
```


Interface	IP Address/VLAN ID	User Name
LS:RI		
ps1.3221225472	192.0.2.10	505024101215074
default:default		
ps1.3221225473	192.0.2.11	505024101215075
default:default		

show subscribers client-type fixed-wireless-access detail (Detail)

user@host> show subscribers client-type fixed-wireless-access detail

```

Type: FWA
User Name: 505024101215074
IP Address: 192.0.2.10
IP Netmask: 255.255.0.0
Interface: ps1.3221225472
Interface type: Dynamic
Dynamic Profile Name: fwa-profile
State: Active
Radius Accounting ID: 1
Session ID: 1
PFE Flow ID: 11
Login Time: 2019-04-10 14:10:12 PDT
Local TEID-U: 1
Local TEID-C: 1
Remote TEID-U: 2000000
Remote TEID-C: 1000000
GTPU Tunnel Remote IP Address: 203.0.113.1.3
GTPU Tunnel Local IP Address: 203.0.113.2.5
GTPC Remote IP Address: 203.0.113.1.2
GTPC Local IP Address: 203.0.113.1.1
Access Point Name: user21

```

show subscribers client-type vlan-oob detail

user@host> show subscribers client-type vlan-oob detail

```

Type: VLAN-OOB
User Name: L2WS.line-aci-1.line-ari-1
Logical System: default
Routing Instance: ISPl
Interface: demux0.1073744127

```

```

Interface type: Dynamic
Underlying Interface: ge-1/0/0
Dynamic Profile Name: Prof_L2WS
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 1234
Session ID: 77
VLAN Id: 126
Core-Facing Interface: ge-2/1/1
VLAN Map Id: 6
Inner VLAN Map Id: 2001
Agent Circuit ID: line-aci-1
Agent Remote ID: line-ari-1
Login Time: 2013-10-29 14:43:52 EDT

```

show subscribers count

```
user@host> show subscribers count
```

```
Total Subscribers: 188, Active Subscribers: 188
```

show subscribers address detail (IPv6)

```
user@host> show subscribers address 203.0.113.137 detail
```

```

Type: PPPoE
User Name: pppoeTerV6User1Svc
IP Address: 203.0.113.137
IP Netmask: 255.0.0.0
IPv6 User Prefix: 2001:db8:0:c88::/32
Logical System: default
Routing Instance: default
Interface: pp0.1073745151
Interface type: Dynamic
Underlying Interface: demux0.8201
Dynamic Profile Name: pppoe-client-profile
MAC Address: 00:00:5e:00:53:53
Session Timeout (seconds): 31622400
Idle Timeout (seconds): 86400
State: Active
Radius Accounting ID: example 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: 203.0.113.29
IP Netmask: 255.255.0.0
Primary DNS Address: 192.0.2.0
Secondary DNS Address: 192.0.2.1
Primary WINS Address: 192.0.2.3
Secondary WINS Address: 192.0.2.4
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:00:5e:00:53:98
State: Active
Radius Accounting ID: example :2304
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: 2001:db8:ffff:1::/32
Logical System: default
Routing Instance: default
Interface: ge-3/1/3.2
Interface type: Static
MAC Address: 00:00:5e:00:53:03

```

```
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
```

show subscribers detail (pseudowire Interface for GRE Tunnel)

user@host> show subscribers detail

Interface	IP Address/VLAN ID	User Name	LS:RI
ps0.3221225484	192.0.2.2		
ps0.3221225485	192.0.2.3		
demux0.3221225486	1		default:default
demux0.3221225487	1		default:default
demux0.3221225488	198.51.0.1		default:default
demux0.3221225489	198.51.0.2		default:default

show subscribers detail (IPv6 Static Demux Interface)

user@host> show subscribers detail

```
Type: STATIC-INTERFACE
User Name: user@example.com
IPv6 Prefix: 2001:db8:3:4:5:6:7:aa/32
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: user@example.com
IP Address: 203.0.113.58
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: user@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: 203.0.113.51
Remote IP Address: 192.0.2.0
Radius Accounting ID: 21
Session ID: 21
Login Time: 2013-01-18 03:01:11 PST
```

```
Type: L2TP
User Name: user@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: 203.0.113.31
Remote IP Address: 192.0.2.1
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@example1.com
IP Address: 203.0.113.13
IPv6 Prefix: 2001:db8:1::/32
IPv6 User Prefix: 2001:db8:1:1::/32
Logical System: default

```

```

Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Dynamic
Dynamic Profile Name: dualStack-Profile1
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: 2
Session ID: 2
Login Time: 2011-11-30 00:18:05 PST

Type: DHCP
IPv6 Prefix: 2001:db8:1::/32
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Static
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: test :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: 203.0.113.15
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: 00:00:5e:00:53:02
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: 2001:db8:ffff:1::/32
Logical System: default
Routing Instance: default
Interface: ge-3/1/3.2
Interface type: Static
MAC Address: 00:00:5e:00:53:03
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
IPv6 Address Pool: pd_pool
IPv6 Network Prefix Length: 48

```

show subscribers extensive (Aggregation Node Interface Set and DSL Forum Attributes)

user@host> show subscribers extensive

```

Type: VLAN-OOB
User Name: ancp
Logical System: default
Routing Instance: ispl-subscriber
Interface: ge-1/0/0.3221225472
Interface type: Dynamic
Interface Set: FRA-DPU-C-100
Underlying Interface: ge-1/0/0
Core IFL Name: ge-1/0/4.0
Dynamic Profile Name: Prof_L2BSA
State: Active
Radius Accounting ID: 1
Session ID: 1
PFE Flow ID: 13
VLAN Id: 50
VLAN Map Id: 20
Inner VLAN Map Id: 1
Inner VLAN Tag Protocol Id: 0x88a8
Agent Circuit ID: circuit 201
Agent Remote ID: remote-id
Aggregation Interface-set Name: FRA-DPU-C-100
Login Time: 2018-05-29 08:43:42 EDT
Accounting interval: 72000
Dynamic configuration:
  junos-cos-scheduler-map: 100m
  junos-inner-vlan-tag-protocol-id: 0x88a8
  junos-vlan-map-id: 20

Type: PPPoE
IP Address: 192.85.128.1
IP Netmask: 255.255.255.255
Logical System: default
Routing Instance: default

```

```

Interface: pp0.3221225474
Interface type: Dynamic
Interface Set: ge-1/0/0
Underlying Interface: demux0.3221225473
Dynamic Profile Name: pppoe-client-profile-with-cos
MAC Address: 00:10:94:00:00:03
State: Active
Radius Accounting ID: 3
Session ID: 3
PFE Flow ID: 16
Stacked VLAN Id: 50
VLAN Id: 7
Agent Circuit ID: circuit 201
Agent Remote ID: remote-id
Aggregation Interface-set Name: FRA-DPU-C-100
Login Time: 2018-05-29 08:43:45 EDT
IP Address Pool: pool-1
Accounting interval: 72000
DSL type: G.fast
Frame/cell mode: Frame
Overhead accounting bytes: 10
Actual upstream data rate: 100000 kbps
Actual downstream data rate: 200000 kbps
Calculated downstream data rate: 180000 kbps
Calculated upstream data rate: 90000 kbps
Adjusted upstream data rate: 80000 kbps
Adjusted downstream data rate: 160000 kbps
DSL Line Attributes
  Agent Circuit ID: circuit 201
  Agent Remote ID: remote-id
  Actual upstream data rate: 100000
  Actual downstream data rate: 200000
  DSL type: G.fast
  Access Aggregation Circuit ID: #FRA-DPU-C-100
  Attribute type: 0xAA, Attribute length: 4
    198 51 100 78

```

show subscribers extensive (Passive Optical Network Circuit Interface Set)

user@host> **show subscribers client-type dhcp extensive**

```

Type: DHCP
IP Address: 192.0.2.136
IP Netmask: 255.255.0.0

```

```

Logical System: default
Routing Instance: default
Interface: demux0.1073741842
Interface type: Dynamic
Interface Set: ot101.xyz101-202
Underlying Interface: demux0.1073741841
Dynamic Profile Name: dhcp-profile
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: user :19
Session ID: 19
VLAN Id: 1100
Agent Remote ID: ABCD01234|100M|AAAA01234|ot101.xyz101-202

Login Time: 2017-03-29 10:30:46 PDT
DHCP Options: len 97
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 02 33 04 00 00
17 70 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 32 2f
32 2d 31 2d 31 37 05 01 06 0f 21 2c 52 2b 02 29 41 42 43 44
30 31 32 33 34 7c 31 30 30 4d 7c 41 41 41 41 30 31 32 33 34
7c 6f 74 6c 30 31 2e 78 79 7a 31 30 31 2d 32 30 32
IP Address Pool: POOL-V4

```

show subscribers extensive (DNS Addresses from Access Profile or Global Configuration)

```
user@host> show subscribers extensive
```

```

Type: DHCP
User Name: test-user@example-com
IP Address: 192.0.2.119
IP Netmask: 255.255.255.255
Domain name server inet: 198.51.100.1 198.51.100.2
IPv6 Address: 2001:db8::1:11
Domain name server inet6: 2001:db8:5001::12 2001:db8:3001::12
Logical System: default
Routing Instance: default
Interface: ge-2/0/3.0
Interface type: Static
Underlying Interface: ge-2/0/3.0
MAC Address: 00:00:5E:00:53:00
State: Active
Radius Accounting ID: 5
Session ID: 5
Login Time: 2017-01-31 11:16:21 IST

```

```

DHCP Options: len 53
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 03 33 04 00 00
00 3c 0c 16 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 35 2f
31 32 2d 30 2d 30 37 05 01 06 0f 21 2c
IP Address Pool: v4-pool

```

show subscribers extensive (DNS Addresses from RADIUS)

```
user@host> show subscribers extensive
```

```

Type: DHCP
User Name: test-user@example-com
IP Address: 192.0.2.119
IP Netmask: 255.255.255.255
Primary DNS Address: 198.51.100.1
Secondary DNS Address: 198.51.100.2
IPv6 Address: 2001:db8::1:11
IPv6 Primary DNS Address: 2001:db8:5001::12
IPv6 Secondary DNS Address: 2001:db8:3001::12
Logical System: default
Routing Instance: default
Interface: ge-2/0/3.0
Interface type: Static
Underlying Interface: ge-2/0/3.0
MAC Address: 00:00:5E:00:53:00
State: Active
Radius Accounting ID: 5
Session ID: 5
Login Time: 2017-01-31 11:16:21 IST
DHCP Options: len 53
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 03 33 04 00 00
00 3c 0c 16 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 35 2f
31 32 2d 30 2d 30 37 05 01 06 0f 21 2c
IP Address Pool: v4-pool

```

show subscribers extensive (IPv4 DNS Addresses from RADIUS, IPv6 from Access Profile or Global Configuration)

```
user@host> show subscribers extensive
```

```

Type: DHCP
User Name: test-user@example-com
IP Address: 192.0.2.119

```

```

IP Netmask: 255.255.255.255
Primary DNS Address: 198.51.100.1
Secondary DNS Address: 198.51.100.2
IPv6 Address: 2001:db8::1:11
Domain name server inet6: 2001:db8:5001::12 2001:db8:3001::12
Logical System: default
Routing Instance: default
Interface: ge-2/0/3.0
Interface type: Static
Underlying Interface: ge-2/0/3.0
MAC Address: 00:00:5E:00:53:00
State: Active
Radius Accounting ID: 5
Session ID: 5
Login Time: 2017-01-31 11:16:21 IST
DHCP Options: len 53
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 03 33 04 00 00
00 3c 0c 16 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 35 2f
31 32 2d 30 2d 30 37 05 01 06 0f 21 2c
IP Address Pool: v4-pool

```

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: user@example.com
IP Address: 203.0.113.58
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@example1.com
IP Address: 203.0.113.13
IPv6 Prefix: 2001:db8:1::/32
IPv6 User Prefix: 2001:db8:1:1::/32
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Dynamic
Dynamic Profile Name: dualStack-Profile1
MAC Address: 00:00:5e:00:53:02

```

```

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: 2001:db8:2016:1:1::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: 2001:db8:1::/32
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Static
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: test :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 (ADF Rules)

```
user@host> show subscribers extensive
```

```

...
Service Session ID: 12
Service Session Name: SERVICE-PROFILE
State: Active
Family: inet
  ADF IPv4 Input Filter Name: __junos_adf_12-demux0.3221225474-inet-in
    Rule 0: 010101000b0101020b020200201811
      from {
        source-address 203.0.113.232;

```

```

        destination-address 198.51.100.0/24;
        protocol 17;
    }
    then {
        accept;
    }

```

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 extensive (PPPoE Subscriber Access Line Rates)

user@host> show subscribers extensive

```

Type: PPPoE
IP Address: 198.51.100.1
IP Netmask: 255.255.255.255
Logical System: default
Routing Instance: default
Interface: pp0.3221225475
Interface type: Dynamic
Underlying Interface: demux0.3221225474
Dynamic Profile Name: pppoe-client-profile-with-cos
MAC Address: 00:00:5e:00:53:02
State: Active
Radius Accounting ID: 4

```


Session ID: 4
 PFE Flow ID: 14
 Stacked VLAN Id: 40
 VLAN Id: 1
 Agent Circuit ID: circuit0
 Agent Remote ID: remote0
 Login Time: 2017-04-06 15:52:32 PDT

User Name: DAVE-L2BSA-SERVICE
 Logical System: default
 Routing Instance: isp-1-subscriber
 Interface: ge-1/2/4.3221225472
 Interface type: Dynamic
 Interface Set: ge-1/2/4
 Underlying Interface: ge-1/2/4
 Core IFL Name: ge-1/3/4.0
 Dynamic Profile Name: L2BSA-88a8-400LL1300VO
 State: Active

Radius Accounting ID: 1

Session ID: 1
 PFE Flow ID: 14
 VLAN Id: 13
 VLAN Map Id: 102
 Inner VLAN Map Id: 1
 Agent Circuit ID: circuit-aci-3
 Agent Remote ID: remote49-3
 Login Time: 2017-04-05 16:59:29 EDT
 Service Sessions: 4
 IFL Input Filter Name: L2BSA-CP-400LL1300VO-ge-1/2/4.3221225472-in
 IFL Output Filter Name: L2BSA-CP-400LL1300VO-ge-1/2/4.3221225472-out
 Accounting interval: 900

DSL type: VDSL

Frame/Cell Mode: Frame

Overhead accounting bytes: -10

Actual upstream data rate: 1024 kbps

Actual downstream data rate: 4096 kbps

Adjusted downstream data rate: 3686 kbps

Adjusted upstream data rate: 922 kbps

Dynamic configuration:

junos-vlan-map-id: 102
 Service Session ID: 5
 Service Session Name: SRL-L1
 State: Active
 Family: inet, inet6

```

IFL Input Filter Name: L2BSA-FWF-in-10048-ge-1/2/4.3221225472-in
IFL Output Filter Name: L2BSA-FWF-out-25088-ge-1/2/4.3221225472-out
Service Activation time: 2017-04-05 16:59:30 EDT
Dynamic configuration:
  l2bsa-fwf-in: L2BSA-FWF-in-10048
  l2bsa-fwf-out: L2BSA-FWF-out-25088
  rldown: 25088
  rlup: 10048

```

show subscribers extensive (Subscriber Session Using PCEF Profile)

user@host> **show subscribers extensive**

```

Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.3221225517
Interface type: Dynamic
Underlying Interface: ge-1/0/3
Dynamic Profile Name: svlan-dhcp
State: Active
Session ID: 59
PFE Flow ID: 71
Stacked VLAN Id: 0x8100.1
VLAN Id: 0x8100.2
Login Time: 2017-03-28 08:23:08 PDT

Type: DHCP
User Name: pcefuser
IP Address: 192.0.2.26
IP Netmask: 255.0.0.0
Logical System: default
Routing Instance: default
Interface: demux0.3221225518
Interface type: Dynamic
Underlying Interface: demux0.3221225517
Dynamic Profile Name: dhcp-client-prof
MAC Address: 00:00:5e:00:53:01
State: Active
Radius Accounting ID: 60
Session ID: 60
PFE Flow ID: 73
Stacked VLAN Id: 1
VLAN Id: 2

```

```

Login Time: 2017-03-28 08:23:08 PDT
Service Sessions: 1
DHCP Options: len 9
35 01 01 37 04 01 03 3a 3b
IP Address Pool: pool-ipv4
IPv4 Input Service Set: tdf-service-set
IPv4 Output Service Set: tdf-service-set
PCEF Profile: pcef-prof-1
PCEF Rule/Rulebase: default
Dynamic configuration:
  junos-input-service-filter: svc-filt-1
  junos-input-service-set: tdf-service-set
  junos-output-service-filter: svc-filt-1
  junos-output-service-set: tdf-service-set
  junos-pcef-profile: pcef-prof-1
  junos-pcef-rule: default

Service Session ID: 61
Service Session Name: pcef-serv-prof
State: Active
Family: inet
IPv4 Input Service Set: tdf-service-set
IPv4 Output Service Set: tdf-service-set
PCEF Profile: pcef-prof-1
PCEF Rule/Rulebase: limit-fb
Service Activation time: 2017-03-28 08:31:19 PDT
Dynamic configuration:
  pcef-prof: pcef-prof-1
  pcef-rule1: limit-fb
  svc-filt: svc-filt-1
  svc-set: tdf-service-set

```

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: 203.0.113.17
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:
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: 203.0.113.17
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: 00:00:5e:00:53:52
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 id accounting-statistics

```
user@host> show subscribers id 601 accounting-statistics
```

```

Session ID: 601
Accounting Statistics:
Input bytes : 199994
Output bytes : 121034
Input packets: 5263
Output packets: 5263
IPv6:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

```

show subscribers interface accounting-statistics

```
user@host> show subscribers interface pp0.3221226949 accounting-statistics
```

```

Session ID: 501
Accounting Statistics:
Input bytes : 199994
Output bytes : 121034
Input packets: 5263
Output packets: 5263
IPv6:

```

```

Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

Session ID: 502
Accounting Statistics:
Input bytes : 87654
Output bytes : 72108
Input packets: 3322
Output packets: 3322
IPv6:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

```

```

Session ID: 503
Accounting Statistics:
Input bytes : 156528
Output bytes : 123865
Input packets: 7448
Output packets: 7448
IPv6:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

```

show subscribers interface extensive

user@host> show subscribers interface demux0.1073741826 extensive

```

Type: VLAN
User Name: user@test.example.com
Logical System: default
Routing Instance: testnet
Interface: demux0.1073741826
Interface type: Dynamic
Dynamic Profile Name: profile-vdemux-relay-23qos
MAC Address: 00:00:5e:00:53:04
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: user@test.example.com
IP Address: 192.0.2.0
IP Netmask: 255.255.255.0
Logical System: default
Routing Instance: testnet
Interface: demux0.1073741826
Interface type: Static
MAC Address: 00:00:5e:00:53:04
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	203.0.113.3	RETAILER1-CLIENT	test1:retailer1
demux0.1073741826	203.0.113.4	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: 203.0.113.37
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: 00:00:5e:00:53:01
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: 203.0.113.2
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: pp0.0
Interface type: Static
MAC Address: 00:00:5e:00:53:02
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 address detail (Enhanced Subscriber Management)

user@host> show subscribers address 203.0.113.111 detail

```

Type: DHCP
User Name: simple_filters_service
IP Address: 203.0.113.111
IP Netmask: 255.0.0.0
Logical System: default

```

```
Routing Instance: default
Interface: demux0.3221225482
Interface type: Dynamic
Underlying Interface: demux0.3221225472
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:00:5e:00:53:0f
State: Active
Radius Accounting ID: 11
Session ID: 11
PFE Flow ID: 15
Stacked VLAN Id: 210
VLAN Id: 209
Login Time: 2014-03-24 12:53:48 PDT
Service Sessions: 1
DHCP Options: len 3
35 01 01
```

show subscribers summary

Syntax

```
show subscribers summary
<all>
<detail | extensive | terse>
<count>
<physical-interface physical-interface-name>
<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

none—Display summary information by state and client type for all subscribers.

all—(Optional) Display summary information by state, client type, and LS:RI.

detail | extensive | terse—(Not supported on MX Series routers) (Optional) Display the specified level of output.

count—(Not supported on MX Series routers) (Optional) Display the count of total subscribers and active subscribers for any specified option.

logical-system *logical-system*—(Optional) Display subscribers whose logical system matches the specified logical system.

physical-interface *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 *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](#) | [849](#)

List of Sample Output

[show subscribers summary on page 899](#)

[show subscribers summary all on page 900](#)

[show subscribers summary physical-interface on page 900](#)

[show subscribers summary physical-interface pic on page 901](#)

[show subscribers summary physical-interface port on page 901](#)

[show subscribers summary physical-interface slot on page 902](#)

[show subscribers summary pic on page 902](#)

[show subscribers summary pic \(Aggregated Ethernet Interfaces\) on page 902](#)

[show subscribers summary port on page 902](#)

[show subscribers summary port \(Pseudowire Interfaces\) on page 903](#)

[show subscribers summary port extensive on page 903](#)

[show subscribers summary slot on page 903](#)

[show subscribers summary terse on page 904](#)

Output Fields

[Table 42 on page 898](#) lists the output fields for the **show subscribers summary** command. Output fields are listed in the approximate order in which they appear.

Table 42: show subscribers summary Output Fields

Field Name	Field Description	Level of Output
Subscribers by State	<p>Number of subscribers summarized by state. The summary information includes the following:</p> <ul style="list-style-type: none"> • 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. 	detail none
Subscribers by Client Type	<p>Number of subscribers summarized by client type. Client types can include DHCP, GRE, L2TP, PPP, PPPoE, STATIC-INTERFACE, VLAN, and VLAN-OOB. Also displays the total number of subscribers for all client types (Total).</p>	detail extensive none
Subscribers by LS:RI	<p>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).</p>	detail none
Subscribers by Connection Type	<p>Number of subscribers summarized by connection type, Cross-connected or Terminated.</p>	extensive
Interface	<p>Interface associated with the subscriber. The router or switch displays subscribers whose interface matches or begins with the specified interface.</p> <p>The * character indicates a continuation of addresses for the same session.</p> <p>For aggregated Ethernet interfaces, the output of the summary (pic port slot) options prefixes the interface name with ae0:.</p> <p>For pseudowire IFDs, this field displays both the pseudowire and the associated logical tunnel (LT) and redundant logical tunnel (RLT) anchor interface. For example:</p> <pre>ps0: lt-2/1/0 ps1: rlt0: lt-4/0/0</pre>	All levels

Table 42: show subscribers summary Output Fields (*continued*)

Field Name	Field Description	Level of Output
Count	<p>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.</p> <p>Multiple pseudowire interfaces can share a given logical tunnel or redundant logical tunnel anchor interface. Starting in Junos OS Release 18.1R1, the field displays subscribers per individual pseudowire interface.</p> <p>In earlier releases, the field displays the same number of subscribers for all pseudowire interfaces that share the same tunnel interface as their anchor point.</p>	detail extensive none
Total Subscribers	Total number of subscribers for all physical interfaces, all PICs, all ports, or all LS:RI slots.	detail extensive none
IP Address/VLAN ID	Subscriber IP address or VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i>	terse
User Name	Name of subscriber.	terse
LS:RI	Logical system and routing instance associated with the subscriber.	terse

Sample Output

show subscribers summary

user@host> **show subscribers summary**

```
Subscribers by State
Init          3
Configured    2
Active       188
Terminating    2
Terminated     1

TOTAL         191
```

Subscribers by Client Type

DHCP	107
PPP	76
VLAN	8
VLAN-OOB	2
TOTAL	196

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-5/0/1       201
ge-5/0/2       301
```

```
Total Subscribers: 502
```

show subscribers summary port (Pseudowire Interfaces)

```
user@host> show subscribers summary port
```

```
ps0: lt-2/1/0 10
ps1: lt-2/1/0 20

Total Subscribers: 30
```

show subscribers summary port extensive

```
user@host>show subscribers summary port extensive
```

```
Interface: ge-5/0/1
Count: 201
Detail:
Subscribers by Client Type
  DHCP: 100
  PPPoE: 100
  VLAN-OOB: 1
Subscribers by Connection Type
  Terminated: 200
  Cross-connected: 1

Interface: ge-5/0/2
Count: 301
Detail:
Subscribers by Client Type
  DHCP: 200
  PPPoE: 100
  VLAN-OOB: 1
Subscribers by Connection Type
  Terminated: 300
  Cross-connected: 1

Total Subscribers: 502
```

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	203.0.113.10	WHOLESALE-CLIENT	default:default
demux0.1073741825	203.0.113.13	RETAILER1-CLIENT	test1:retailer1
demux0.1073741826	203.0.113.213	RETAILER2-CLIENT	test1:retailer2

show vpls connections

Syntax

```
show vpls connections
<brief | extensive>
<down | up | up-down>
<history>
<instance instance-name local-site local-site-name remote-site remote-site-name>
<instance-history>
<logical-system (all | logical-system-name)>
<status>
<summary>
```

Release Information

Command introduced before Junos OS Release 7.4.

instance-history option introduced in Junos OS Release 12.3R2.

Description

(T Series and M Series routers, except for the M160 router) Display virtual private LAN service (VPLS) connection information.

Options

none—Display information about all VPLS connections for all routing instances.

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

down | up | up-down—(Optional) Display nonoperational, operational, or both types of connections.

history—(Optional) Display information about connection history.

instance *instance-name*—(Optional) Display the VPLS connections for the specified routing instance only.

instance-history—(Optional) Display information about connection history for a particular instance.

local-site *local-site-name*—(Optional) Display the VPLS connections for the specified local site name or ID only.

remote-site *remote-site-name*—(Optional) Display the VPLS connections for the specified remote site name or ID only. Label block size information is always shown as 0 when using this option.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

status—(Optional) Display information about the connection and interface status.

summary—(Optional) Display summary of all VPLS connections information.

Required Privilege Level

view

List of Sample Output[show vpls connections on page 914](#)[show vpls connections \(with FEC128 and FEC129 in the same routing-instance\) on page 916](#)[show vpls connections \(with multiple pseudowires\) on page 917](#)[show vpls connections extensive \(Static VPLS Neighbors\) on page 918](#)**Output Fields**

[Table 43 on page 906](#) lists the output fields for the **show vpls connections** command. Output fields are listed in the approximate order in which they appear.

Table 43: show vpls connections Output Fields

Field Name	Field Description
Instance	Name of the VPLS instance.
Local site	Name of the local site.
VPLS-id	Identifier for the VPLS site.
Number of local interfaces	Number of interfaces configured for the local site.
Number of local interfaces up	Number of interfaces configured for the local site that are currently up.
IRB interface present	Indicates whether or not an integrated routing and bridging (IRB) interface is present (yes or no).

Table 43: show vpls connections Output Fields (*continued*)

Field Name	Field Description
Intf	<p>List of all of the interfaces configured for the local site. The types of interfaces can include VPLS virtual loopback tunnel interfaces and label-switched interfaces. Any interface that supports VPLS could be listed here.</p> <p>Virtual loopback tunnel interfaces are displayed using the vt-fpc/pic/port.nnnnn format. Label-switched interfaces are displayed using the lsi.nnnnn format. In both cases, nnnnn is a dynamically generated virtual port used to transport and receive packets from other provider edge (PE) routers in the VPLS domain.</p> <p>Each interface might include the following information:</p> <ul style="list-style-type: none"> • Identification as a VPLS interface • Name of the associated VPLS routing instance • Local site number • Remote site number • VPLS neighbor address • VPLS identifier
Interface flags	<p>Flag associated with the interface. Can include the following:</p> <ul style="list-style-type: none"> • VC-Down—The virtual circuit associated with this interface is down.
Label-base	First label in a block of labels. A remote PE router uses this first label when sending traffic toward the advertising PE router.
Offset	Displays the VPLS Edge (VE) block offset in the Layer 2 VPN NLRI. The VE block offset is used to identify a label block from which a particular label value is selected to setup a pseudowire for a remote site. The block offset value itself indicates the starting VE ID that maps to the label base contained in the VPLS NLRI advertisement.
Size	<p>Label block size. A configurable value that represents the number of label blocks required to cover all the pseudowires for the remote peer.</p> <p>Acceptable configuration values are: 2, 4, 8 and 16. The default value is 2. A value of 0 will be displayed when using the remote-site option.</p>
Range	Label block range. A value that keeps track of the numbers of remote sites discovered within each label block.

Table 43: show vpls connections Output Fields (*continued*)

Field Name	Field Description
Preference	Preference value advertised for a VPLS site. When multiple PE routers are assigned the same VE ID for multihoming, you might need to specify that a particular PE router acts as the designated forwarder by configuring the site preference value. The site preference indicates the degree of preference for a particular customer site. The site preference is one of the tie-breaking criteria used in a designated forwarder election.
status-vector	Bit vector advertising the state of local PE-CE circuits to remote PE routers. A bit value of 0 indicates that the local circuit and LSP tunnel to the remote PE router are up, whereas a value of 1 indicates either one or both are down.
connection-site	Name of the connection site.
Neighbor	IP address and VPLS identifier for the VPLS neighbor. If multiple pseudowires have been configured, the IP address will also show the PW-specific <i>vpls-id-list</i> , for example, 203.0.113.144 (vpls-id 200).
Type	Type of connection: loc (local) or rmt (remote).

Table 43: show vpls connections Output Fields (continued)

Field Name	Field Description
St	

Table 43: show vpls connections Output Fields (*continued*)

Field Name	Field Description
	<p>Status of the VPLS connection (corresponds with Legend for Connection Status):</p> <ul style="list-style-type: none"> • EI—The local VPLS interface is configured with an encapsulation that is not supported. • EM—The encapsulation type received on this VPLS connection from the neighbor does not match the local VPLS connection interface encapsulation type. • VC-Dn—The virtual circuit is currently down. • CM—The two routers do not agree on a control word, which causes a control word mismatch. • CN—The virtual circuit is not provisioned properly. • OR—The label associated with the virtual circuit is out of range. • OL—No advertisement has been received for this virtual circuit from the neighbor. There is no outgoing label available for use by this virtual circuit. • LD—All of the CE-facing interfaces to the local site are down. Therefore, the connection to the local site is signaled as down to the other PE routers. No pseudowires can be established. • RD—All the interfaces to the remote neighbor are down. Therefore, the remote site has been signaled as down to the other PE routers. No pseudowires can be established. • LN—The local site has lost path selection to the remote site and therefore no pseudowires can be established from this local site. • RN—The remote site has lost path selection to a local site or other remote site and therefore no pseudowires are established to this remote site. <p>In a multihoming configuration, one multihomed PE site displays the state LN, and the other multihomed PE site displays the state RN in the following circumstances:</p> <ul style="list-style-type: none"> • The multihomed links are both configured to be the backup site. • The two multihomed PE routers have the same site ID, but have a peering relationship with a route reflector (RR) that has a different site ID. <ul style="list-style-type: none"> • XX—The VPLS connection is down for an unknown reason. This is a programming error. • MM—The MTU for the local site and the remote site do not match. • BK—The router is using a backup connection.

Table 43: show vpls connections Output Fields (continued)

Field Name	Field Description
	<ul style="list-style-type: none">• PF—Profile parse failure.• RS—The remote site is in a standby state.• NC—The interface encapsulation is not configured as an appropriate CCC, TCC, or VPLS encapsulation.• WE—The encapsulation configured for the interface does not match the encapsulation configured for the associated connection within the VPLS routing instance.

Table 43: show vpls connections Output Fields (continued)

Field Name	Field Description
	<ul style="list-style-type: none"> • NP—The router detects that interface hardware is not present. The hardware might be offline, a PIC might not be of the desired type, or the interface might be configured in a different routing instance. • -->—Only the outbound connection is up. • <--—Only the inbound connection is up. • Up—The VPLS connection is operational. • Dn—The VPLS connection is down. • CF—The router cannot find enough bandwidth to the remote router to satisfy the VPLS connection bandwidth requirement. • SC—The local site identifier matches the remote site identifier. No pseudowire can be established between these two sites. You should configure different values for the local and remote site identifiers. • LM—The local site identifier is not the minimum designated, meaning it is not the lowest. There is another local site with a lower site identifier. Pseudowires are not being established to this local site. and the associated local site identifier is not being used to distribute VPLS label blocks. However, this is not an error state. Traffic continues to be forwarded to the PE router interfaces connected to the local sites when the local sites are in this state. • RM—The remote site identifier is not the minimum designated, meaning it is not the lowest. There is another remote site connected to the same PE router which has lower site identifier. The PE router cannot established a pseudowire to this remote site and the associated remote site identifier cannot be used to distribute VPLS label blocks. However, this is not an error state. Traffic can continue to be forwarded to the PE router interface connected to this remote site when the remote site is in this state. • IL—The incoming packets for the VPLS connection have no MPLS label. • MI—The configured mesh group identifier is in use by another system in the network. • ST—The router has switched to a standby connection. • PB—Profile busy. • SN—The VPLS neighbor is static.
Time last up	Time connection was last in the Up condition.
# Up trans	Number of transitions from Down to Up condition.

Table 43: show vpls connections Output Fields (*continued*)

Field Name	Field Description
Status	<p>Status of the (local or remote circuit) local interface:</p> <ul style="list-style-type: none"> • Up—Operational • Dn—Down • NP—Not present • DS—Disabled • WE—Wrong encapsulation • UN—Uninitialized
Encapsulation	Type of encapsulation: VPLS .
Remote PE	Address of the remote provider edge router.
Negotiated control-word	Whether a control word has been negotiated: Yes or No .
Incoming label	Name of the incoming label.
Outgoing label	Name of the outgoing label.
Negotiated PW status TLV	Indicates whether or not the pseudowire status TLV has been negotiated for the VPLS connection.
Local interface	<p>Provides the following information about the local interface configured for the VPLS neighbor:</p> <ul style="list-style-type: none"> • Name of the local interface • Status—Interface status (Up or Down) • Encapsulation—Interface encapsulation (for example, ETHERNET) • Description—Includes the VPLS instance name, the VPLS neighbor address, and the VPLS identifier
Time	Date and time of VPLS connection event.
Event	Type of event.
Interface/Lbl/PE	Interface, label, or PE router.

Table 43: show vpls connections Output Fields (*continued*)

Field Name	Field Description
Connection History	Each entry can include the date, time, year, and the connection event. Connection events include any of a variety of events related to VPLS connections, such as route changes, label updates, and interfaces going down or coming up.

Sample Output

show vpls connections

user@host> show vpls connections

Layer-2 VPN connections:

Legend for connection status (St)

EI -- encapsulation invalid	NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch	WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down	NP -- interface hardware not present
CM -- control-word mismatch	-< -- only outbound connection is up
CN -- circuit not provisioned	>- -- only inbound connection is up
OR -- out of range	Up -- operational
OL -- no outgoing label	Dn -- down
LD -- local site signaled down	CF -- call admission control failure
RD -- remote site signaled down	SC -- local and remote site ID collision
LN -- local site not designated	LM -- local site ID not minimum designated
RN -- remote site not designated	RM -- remote site ID not minimum designated
XX -- unn connection status	IL -- no incoming label
MM -- MTU mismatch	MI -- Mesh-Group ID not available
BK -- Backup connection	ST -- Standby connection
PF -- Profile parse failure	PB -- Profile busy

Legend for interface status

Up -- operational
Dn -- down

Instance: vpls-1

Local site: 1 (11)

Number of local interfaces: 1

Number of local interfaces up: 1

IRB interface present: no

```

lt-1/3/0.10496
vt-1/3/0.1048588    1          Intf - vpls vpls-1 local site 11 remote site 1
vt-1/2/0.1048591    2          Intf - vpls vpls-1 local site 11 remote site 2
vt-1/2/0.1048585    3          Intf - vpls vpls-1 local site 11 remote site 3
vt-1/2/0.1048587    4          Intf - vpls vpls-1 local site 11 remote site 4
vt-1/2/0.1048589    5          Intf - vpls vpls-1 local site 11 remote site 5
vt-1/3/0.1048586    6          Intf - vpls vpls-1 local site 11 remote site 6
vt-1/3/0.1048590    7          Intf - vpls vpls-1 local site 11 remote site 7
vt-1/3/0.1048584    8          Intf - vpls vpls-1 local site 11 remote site 8

```

Label-base	Offset	Size	Range	Preference
800256	1	16	16	100

Timer Values:

```

Startup wait time: 120 seconds
New site wait-time: 20 seconds
Collision detect time: 30 seconds
Reclaim wait time: 748 milliseconds

```

connection-site	Type	St	Time last up	# Up trans
1	rmt	Up	Apr 28 13:28:24 2009	2

Remote PE: 192.0.2.1, Negotiated control-word: No

Incoming label: 800256, Outgoing label: 800026

Local interface: vt-1/3/0.1048588, Status: Up, Encapsulation: VPLS

Description: Intf - vpls vpls-1 local site 11 remote site 1

Connection History:

```

Apr 28 13:28:24 2009 status update timer
Apr 28 13:28:24 2009 PE route down
Apr 28 13:24:27 2009 status update timer
Apr 28 13:24:27 2009 loc intf up vt-1/3/0.1048588
Apr 28 13:24:27 2009 PE route changed
Apr 28 13:24:27 2009 Out lbl Update 800026
Apr 28 13:24:27 2009 In lbl Update 800256
Apr 28 13:24:27 2009 loc intf down

```

2	rmt	Up	Apr 28 13:28:24 2009	2
---	-----	----	----------------------	---

Remote PE: 192.0.2.71, Negotiated control-word: No

Incoming label: 800257, Outgoing label: 800034

Local interface: vt-1/2/0.1048591, Status: Up, Encapsulation: VPLS

Description: Intf - vpls vpls-1 local site 11 remote site 2

Connection History:

```

Apr 28 13:28:24 2009 status update timer
Apr 28 13:28:24 2009 PE route down
Apr 28 13:24:28 2009 status update timer
Apr 28 13:24:28 2009 loc intf up vt-1/2/0.1048591
Apr 28 13:24:28 2009 PE route changed
Apr 28 13:24:28 2009 Out lbl Update 800034

```

```

Apr 28 13:24:28 2009 In lbl Update      800257
Apr 28 13:24:28 2009 loc intf down

```

show vpls connections (with FEC128 and FEC129 in the same routing-instance)

user@host> show vpls connections

```

Instance: fec129
  L2vpn-id: 1:1
  Local-id: 203.0.113.0
FEC129-VPLS State:
  Mesh-group connections: __ves__
    Remote-id      Type  St      Time last up      # Up trans
    203.0.3.3      rmt   Up      Sep 19 09:59:56 2017      1
    Remote PE: 203.0.3.3, Negotiated control-word: No
    Incoming label: 262155, Outgoing label: 262164
    Negotiated PW status TLV: No
    Local interface: lsi.1048844, Status: Up, Encapsulation: ETHERNET
    Description: Intf - vpls fec129 local-id 10.4.4.4 remote-id 203.0.3.3
neighbor 203.0.3.3
  Flow Label Transmit: No, Flow Label Receive: No
    203.0.2.2      rmt   Up      Sep 19 09:59:52 2017      1
    Remote PE: 203.0.2.2, Negotiated control-word: No
    Incoming label: 262154, Outgoing label: 262157
    Negotiated PW status TLV: No
    Local interface: lsi.1048846, Status: Up, Encapsulation: ETHERNET
    Description: Intf - vpls fec129 local-id 10.4.4.4 remote-id 203.0.2.2
neighbor 203.0.2.2
  Flow Label Transmit: No, Flow Label Receive: No
    203.0.1.1      rmt   Up      Sep 19 09:59:48 2017      1
    Remote PE: 203.0.1.1, Negotiated control-word: No
    Incoming label: 262156, Outgoing label: 262157
    Negotiated PW status TLV: No
    Local interface: lsi.1048845, Status: Up, Encapsulation: ETHERNET
    Description: Intf - vpls fec129 local-id 10.4.4.4 remote-id 203.0.1.1
neighbor 203.0.1.1
  Flow Label Transmit: No, Flow Label Receive: No

LDP-VPLS State
  Mesh-group connections: MG1
    Neighbor      Type  St      Time last up      # Up trans
    203.0.6.6(vpls-id 1)  rmt   Up      Sep 17 19:17:11 2017      1
    Remote PE: 203.0.6.6, Negotiated control-word: No
    Incoming label: 262423, Outgoing label: 262145

```



```

Negotiated PW status TLV: No
Local interface: lsi.1049859, Status: Up, Encapsulation: ETHERNET
  Description: Intf - vpls bgp-vpls neighbor 203.0.6.6 vpls-id 1
Flow Label Transmit: No, Flow Label Receive: No
203.0.7.7(vpls-id 1)      rmt    Up      Sep 17 19:17:04 2017      1
Remote PE: 203.0.7.7, Negotiated control-word: No
Incoming label: 262424, Outgoing label: 262145
Negotiated PW status TLV: No
Local interface: lsi.1049857, Status: Up, Encapsulation: ETHERNET
  Description: Intf - vpls bgp-vpls neighbor 203.0.7.7 vpls-id 1
Flow Label Transmit: No, Flow Label Receive: No
Mesh-group connections: MG2
Neighbor                  Type  St      Time last up          # Up trans
203.0.5.5(vpls-id 1)      rmt    Up      Sep 17 19:17:00 2017      1
Remote PE: 203.0.5.5, Negotiated control-word: No
Incoming label: 262425, Outgoing label: 299872
Negotiated PW status TLV: No
Local interface: lsi.1049856, Status: Up, Encapsulation: ETHERNET
  Description: Intf - vpls bgp-vpls neighbor 203.0.5.5 vpls-id 1
Flow Label Transmit: No, Flow Label Receive: No

```

show vpls connections (with multiple pseudowires)

user@host> show vpls connections

Layer-2 VPN connections:

Legend for connection status (St)

EI -- encapsulation invalid	NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch	WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down	NP -- interface hardware not present
CM -- control-word mismatch	-> -- only outbound connection is up
CN -- circuit not provisioned	<- -- only inbound connection is up
OR -- out of range	Up -- operational
OL -- no outgoing label	Dn -- down
LD -- local site signaled down	CF -- call admission control failure
RD -- remote site signaled down	SC -- local and remote site ID collision
LN -- local site not designated	LM -- local site ID not minimum designated
RN -- remote site not designated	RM -- remote site ID not minimum designated
XX -- unknown connection status	IL -- no incoming label
MM -- MTU mismatch	MI -- Mesh-Group ID not available
BK -- Backup connection	ST -- Standby connection
PF -- Profile parse failure	PB -- Profile busy
RS -- remote site standby	SN -- Static Neighbor

```

LB -- Local site not best-site   RB -- Remote site not best-site
VM -- VLAN ID mismatch

Legend for interface status
Up -- operational
Dn -- down

Instance: vpls
  VPLS-id: 100
  Mesh-group connections: __ves__
    Neighbor                Type  St    Time last up          # Up trans
    10.255.114.3 (vpls-id 100) rmt   Up    Apr 11 23:38:38 2013          1
    Remote PE: 10.255.114.3, Negotiated control-word: No
    Incoming label: 262145, Outgoing label: 262145
    Negotiated PW status TLV: No
    Local interface: lsi.1049090, Status: Up, Encapsulation: ETHERNET
    Description: Intf - vpls h-vpls neighbor 10.255.114.3 vpls-id 100

  Mesh-group connections: spokes
    Neighbor                Type  St    Time last up          # Up trans
    10.255.114.4 (vpls-id 200) rmt   Up    Apr 11 23:39:25 2013          1
    Remote PE: 10.255.114.4, Negotiated control-word: No
    Incoming label: 262148, Outgoing label: 304224
    Negotiated PW status TLV: Yes
    local PW status code: 0x00000000, Neighbor PW status code: 0x00000000
    Local interface: lsi.1049091, Status: Up, Encapsulation: ETHERNET
    Description: Intf - vpls h-vpls neighbor 10.255.114.4 vpls-id 200
    10.255.114.4 (vpls-id 201) rmt   Up    Apr 11 23:39:25 2013          1
    Remote PE: 10.255.114.4, Negotiated control-word: No
    Incoming label: 262149, Outgoing label: 304225
    Negotiated PW status TLV: Yes
    local PW status code: 0x00000000, Neighbor PW status code: 0x00000000
    Local interface: lsi.1049096, Status: Up, Encapsulation: ETHERNET
    Description: Intf - vpls h-vpls neighbor 10.255.114.4 vpls-id 201

```

show vpls connections extensive (Static VPLS Neighbors)

user@host> show vpls connections extensive instance red

```

Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same

```

```

VC-Dn -- Virtual circuit down      NP -- interface hardware not present
CM -- control-word mismatch        -> -- only outbound connection is up
CN -- circuit not provisioned      <- -- only inbound connection is up
OR -- out of range                 Up -- operational
OL -- no outgoing label            Dn -- down
LD -- local site signaled down     CF -- call admission control failure
RD -- remote site signaled down    SC -- local and remote site ID collision
LN -- local site not designated    LM -- local site ID not minimum designated
RN -- remote site not designated   RM -- remote site ID not minimum designated
XX -- unn connection status        IL -- no incoming label
MM -- MTU mismatch                 MI -- Mesh-Group ID not available
BK -- Backup connection            ST -- Standby connection
PF -- Profile parse failure        PB -- Profile busy
RS -- remote site standby          SN -- Static Neighbor

```

Legend for interface status

```

Up -- operational
Dn -- down

```

Instance: static

VPLS-id: 1

Number of local interfaces: 1

Number of local interfaces up: 1

ge-0/0/5.0

lsi.1049344 Intf - vpls static neighbor 10.255.114.3 vpls-id

1

Neighbor	Type	St	Time last up	# Up trans
10.255.114.3(vpls-id 1)(SN)	rmt	Up	Mar 4 08:48:41 2010	1

Remote PE: 10.255.114.3, Negotiated control-word: No

Incoming label: 29696, Outgoing label: 29697

Negotiated PW status TLV: No

Local interface: lsi.1049344, Status: Up, Encapsulation: ETHERNET

Description: Intf - vpls static neighbor 10.255.114.3 vpls-id 1

Connection History:

Mar 4 08:48:41 2010	status update timer	
Mar 4 08:48:41 2010	PE route changed	
Mar 4 08:48:41 2010	Out lbl Update	29697
Mar 4 08:48:41 2010	In lbl Update	29696
Mar 4 08:48:41 2010	loc intf up	lsi.1049344

user@PE1> **show vpls connections extensive** (Multihoming with FEC 129)

Layer-2 VPN connections:

Legend for connection status (St)

EI -- encapsulation invalid	NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch	WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down	NP -- interface hardware not present
CM -- control-word mismatch	-> -- only outbound connection is up
CN -- circuit not provisioned	<- -- only inbound connection is up
OR -- out of range	Up -- operational
OL -- no outgoing label	Dn -- down
LD -- local site signaled down	CF -- call admission control failure
RD -- remote site signaled down	SC -- local and remote site ID collision
LN -- local site not designated	LM -- local site ID not minimum designated
RN -- remote site not designated	RM -- remote site ID not minimum designated
XX -- unknown connection status	IL -- no incoming label
MM -- MTU mismatch	MI -- Mesh-Group ID not available
BK -- Backup connection	ST -- Standby connection
PF -- Profile parse failure	PB -- Profile busy
RS -- remote site standby	SN -- Static Neighbor
LB -- Local site not best-site	RB -- Remote site not best-site
VM -- VLAN ID mismatch	

Legend for interface status

Up -- operational

Dn -- down

Instance: green

L2vpn-id: 100:100

Local-id: 192.0.2.2

Number of local interfaces: 2

Number of local interfaces up: 2

ge-0/3/1.0

ge-0/3/3.0

lsi.101711873

Intf - vpls green local-id 192.0.2.2 remote-id

192.0.2.4 neighbor 192.0.2.4

Remote-id	Type	St	Time last up	# Up trans
192.0.2.4	rmt	Up	Jan 31 13:49:52 2012	1

Remote PE: 192.0.2.4, Negotiated control-word: No

Incoming label: 262146, Outgoing label: 262146

Local interface: lsi.101711873, Status: Up, Encapsulation: ETHERNET

Description: Intf - vpls green local-id 192.0.2.2 remote-id 192.0.2.4

neighbor 192.0.2.4

Connection History:

Jan 31 13:49:52 2012 status update timer

```

Jan 31 13:49:52 2012 PE route changed
Jan 31 13:49:52 2012 Out lbl Update                262146
Jan 31 13:49:52 2012 In lbl Update                262146
Jan 31 13:49:52 2012 loc intf up                  lsi.101711873
Multi-home:
Local-site      Id      Pref   State
test            1      100    Up
Number of interfaces: 1
Number of interfaces up: 1
ge-0/3/1.0
Received multi-homing advertisements:
Remote-PE      Pref   flag   Description
192.0.2.4      100    0x0

```

show vpls flood event-queue

Syntax

```
show vpls flood event-queue
```

Release Information

Command introduced in Junos OS Release 8.0.

Description

Display the pending events in the VPLS flood queue.

Options

This command has no options.

Required Privilege Level

view

List of Sample Output

[show vpls flood event-queue on page 923](#)

Output Fields

[Table 44 on page 922](#) lists the output fields for the **show vpls flood event-queue** command. Output fields are listed in the approximate order in which they appear.

Table 44: show vpls flood event-queue Output Fields

Field Name	Field Description
Current Pending Event	Provides information on the current event in the VPLS flood event queue.
Name	Name of the event.
Owner Name	Name of the interface associated with the flood event.
Pending Op	Pending operation for the event.
Last Error	Name of the last error encountered.
Number of Retries	Number of attempts made to update the event queue.
Pending Event List	List of the events awaiting processing.
Event Name	Name of the event.

Table 44: show vpls flood event-queue Output Fields (*continued*)

Field Name	Field Description
Pending Op	Pending operation for the event.
Event Identifier	Name of the interface associated with the flood event.

Sample Output

show vpls flood event-queue

user@host> **show vpls flood event-queue**

Current Pending Event

Name: Flood Nexthop

Owner Name:ge-4/3/0.0

Pending Op: ADD

Last Error:ENOMEM

Number of Retries:3

Pending Event List:

Event Name	Pending Op	Event Identifier
Flood Nexthop	ADD	ge-4/3/0.0
Flood Route	ADD	ge-4/3/0.0

show vpls flood instance

Syntax

```
show vpls flood instance
<brief | detail | extensive>
<instance-name>
<logical-system logical-system-name>
```

Release Information

Command introduced in Junos OS Release 8.0.

Description

Display VPLS information related to the flood process.

Options

none—Display VPLS information related to the flood process for all routing instances.

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

instance-name—(Optional) Display VPLS information related to the flood process for the specified routing instance.

logical-system logical-system-name—(Optional) Display VPLS information related to the flood process for the specified logical system.

Required Privilege Level

view

List of Sample Output

[show vpls flood instance on page 925](#)

[show vpls flood instance logical-system-name on page 925](#)

[show vpls flood instance detail on page 926](#)

Output Fields

[Table 45 on page 924](#) lists the output fields for the **show vpls flood instance** command. Output fields are listed in the approximate order in which they appear.

Table 45: show vpls flood instance Output Fields

Field Name	Field Description
Logical system	Name of the logical system.
Name	Name of the VPLS routing instance.

Table 45: show vpls flood instance Output Fields (*continued*)

Field Name	Field Description
CEs	Number of CE routers connected to the VPLS instance.
VEs	Number of VE routers connected to the VPLS instance.
Flood routes	List of all flood routes associated with the VPLS instance.
Prefix	Prefix for the route.
Type	Type of route.
Owner	VPLS routing instance or interface associated with the route.
NhType	Next-hop type. For example, flood for a flood route.
Nhindex	Next-hop index number for the route.

Sample Output

show vpls flood instance

```
user@host> show vpls flood instance
```

```
Logical system: __example_ls1__
Name: green
CEs: 1
VEs: 1
Flood Routes:
  Prefix    Type           Owner           NhType    NhIndex
  default   ALL_CE_FLOOD   green           flood      383
  0x47/16   CE_FLOOD       fe-1/2/1.0      flood      388
```

show vpls flood instance logical-system-name

```
user@host:__example_ls1__> show vpls flood instance example_ls1
```

```
Logical system: __example_ls1__
Name: green
CEs: 1
VEs: 1
Flood Routes:
  Prefix    Type           Owner           NhType      NhIndex
  default   ALL_CE_FLOOD  green           flood        383
  0x47/16   CE_FLOOD      fe-1/2/1.0     flood        388
```

show vpls flood instance detail

user@host:__example_ls1__> **show vpls flood instance detail**

```
Logical system: __example_ls1__
Name: green
CEs: 1
VEs: 1
Flood Routes:
  Prefix    Type           Owner           NhType      NhIndex
  default   ALL_CE_FLOOD  green           flood        383
  0x47/16   CE_FLOOD      fe-1/2/1.0     flood        388
```

show vpls flood route

Syntax

```
show vpls flood route
(all-ce-flood instance-name instance-name <logical-system-name logical-system-name> |
ce-flood interface interface-name)
```

Release Information

Command introduced in Junos OS Release 8.0.

Description

Display VPLS route information related to the flood process for either the specified routing instance or the specified interface.

Options

all-ce-flood—Display the flood next-hop route for all customer edge routers for traffic coming from the core of the network.

ce-flood interface *interface-name*—Display the flood next-hop route for traffic coming from the specified customer edge interface.

instance-name *instance-name*—Display the flood routes for the specified instance.

logical-system-name *logical-system-name*—(Optional) Specify the logical system whose flood routes you want to display. You can only specify the default logical system name for VPLS. The default logical system name is `__example_ls1__` (the name must be entered in the command with the underscore characters).

Required Privilege Level

view

List of Sample Output

[show vpls flood route all-ce-flood on page 928](#)

[show vpls flood route ce-flood on page 929](#)

Output Fields

[Table 46 on page 927](#) lists the output for the **show vpls flood route** command. Output fields are listed in the approximate order in which they appear.

Table 46: show vpls flood route Output Fields

Field Name	Field Description
Flood route prefix	Prefix for the flood route.

Table 46: show vpls flood route Output Fields (continued)

Field Name	Field Description
Flood route type	Type of flood route (either CE_FLOOD or ALL_CE_FLOOD).
Flood route owner	VPLS routing instance or interface associated with the flood route.
Nexthop type	Next-hop type. For example, flood for a flood route.
Nexthop index	Next-hop index number for the route.
Interfaces flooding to	Interfaces to which VPLS routes are being flooded.
Name	Name of the interface.
Type	Type of VPLS router (CE or VE).
Nh type	Next-hop type.
Index	Index number for the flood route.

Sample Output

show vpls flood route all-ce-flood

```
user@host:__example_ls1__> show vpls flood route all-ce-flood logical-system-name
__example_ls1__instance-name green
```

```
Flood route prefix: default
Flood route type: ALL_CE_FLOOD
Flood route owner: green
Nexthop type: flood
Nexthop index: 383
  Interfaces Flooding to:
  Name                Type          NhType      Index
  fe-1/2/1.0          CE
```

show vpls flood route ce-flood

user@host: __example_ls1__> **show vpls flood route ce-flood interface fe-1/2/1.0**

```
Flood route prefix: 0x47/16
Flood route type: CE_FLOOD
Flood route owner: fe-1/2/1.0
Nexthop type: flood
Nexthop index: 388
  Interfaces Flooding to:
  Name                Type      NhType      Index
  lsi.49152            VE        indr        262142
```

show vpls mac-table

Syntax

```
show vpls mac-table
<age>
<brief | detail | extensive | summary>
<bridge-domain bridge-domain-name>
<instance instance-name>
<interface interface-name>
<logical-system (all | logical-system-name)>
<mac-address>
<vlan-id vlan-id-number>
```

Release Information

Command introduced in Junos OS Release 8.5.

Command introduced in Junos OS Release 15.1.

Description

Display learned virtual private LAN service (VPLS) media access control (MAC) address information.

Options

none—Display all learned VPLS MAC address information.

age— (Optional) Display age of a single mac-address.

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

bridge-domain *bridge-domain-name*—(Optional) Display learned VPLS MAC addresses for the specified bridge domain.

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

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

logical-system (all | *logical-system-name*)—(Optional) Display learned VPLS MAC addresses for all logical systems or for the specified logical system.

mac-address—(Optional) Display the specified learned VPLS MAC address information..

vlan-id *vlan-id-number*—(Optional) Display learned VPLS MAC addresses for the specified VLAN.

Required Privilege Level

view

List of Sample Output

[show vpls mac-table on page 932](#)

[show vpls mac-table \(with Layer 2 Services over GRE Interfaces\) on page 932](#)

[show vpls mac-table \(with VXLAN enabled\) on page 933](#)

[show vpls mac-table age \(for GE interface\) on page 933](#)

[show vpls mac-table age \(for AE interface\) on page 933](#)

[show vpls mac-table count on page 934](#)

[show vpls mac-table detail on page 935](#)

[show vpls mac-table extensive on page 935](#)

Output Fields

[Table 47 on page 931](#) describes the output fields for the **show vpls mac-table** command. Output fields are listed in the approximate order in which they appear.

Table 47: show vpls mac-table Output fields

Field Name	Field Description
Age	Age of a single mac-address.
Routing instance	Name of the routing instance.
Bridging domain	Name of the bridging domain.
MAC address	MAC address or addresses learned on a logical interface.
MAC flags	Status of MAC address learning properties for each interface: <ul style="list-style-type: none"> • S—Static MAC address configured. • D—Dynamic MAC address learned. • SE—MAC accounting is enabled. • NM—Nonconfigured MAC.
Logical interface	Name of the logical interface.
MAC count	Number of MAC addresses learned on a specific routing instance or interface.
Learning interface	Logical interface or logical Label Switched Interface (LSI) the address is learned on.
Base learning interface	Base learning interface of the MAC address. This field is introduced in Junos OS Release 14.2.
Learn VLAN ID/VLAN	VLAN ID of the routing instance or bridge domain in which the MAC address was learned.
VXLAN ID/VXLAN	VXLAN Network Identifier (VNI)

Table 47: show vpls mac-table Output fields (*continued*)

Field Name	Field Description
Layer 2 flags	Debugging flags signifying that the MAC address is present in various lists.
Epoch	Spanning Tree Protocol epoch number identifying when the MAC address was learned. Used for debugging.
Sequence number	Sequence number assigned to this MAC address. Used for debugging.
Learning mask	Mask of Packet Forwarding Engines where this MAC address was learned. Used for debugging.
IPC generation	Creation time of the logical interface when this MAC address was learned. Used for debugging.

Sample Output

show vpls mac-table

```
user@host> show vpls mac-table
```

```
MAC flags (S -static MAC, D -dynamic MAC,
           SE -Statistics enabled, NM -Non configured MAC)
```

```
Routing instance : vpls_ldp1
```

```
VLAN : 223
```

MAC address	MAC flags	Logical interface
00:00:5e:00:53:5d	D	ge-0/2/5.400

```
MAC flags (S -static MAC, D -dynamic MAC,
           SE -Statistics enabled, NM -Non configured MAC)
```

```
Routing instance : vpls_red
```

```
VLAN : 401
```

MAC address	MAC flags	Logical interface
00:00:5e:00:53:12	D	lsi.1051138
00:00:5e:00:53:f0	D	lsi.1051138

show vpls mac-table (with Layer 2 Services over GRE Interfaces)

```
user@host> show vpls mac-table
```



```
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned
          SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : vpls_4site:1000
Bridging domain : __vpls_4site:1000__,   MAC          MAC          Logical
address          flags          interface
00:00:5e:00:53:f4  D,SE          ge-4/2/0.1000
00:00:5e:00:53:33  D,SE          lsi.1052004
00:00:5e:00:53:32  D,SE          lsi.1048840
00:00:5e:00:53:14  D,SE          lsi.1052005
00:00:5e:00:53:f7  D,SE          gr-1/2/10.10
```

show vpls mac-table (with VXLAN enabled)

```
user@host> show vpls mac-table
```

```
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned
          SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : vpls_4site:1000
Bridging domain : __vpls_4site:1000__, VLAN : 4094,4093
VXLAN: Id : 300, Multicast group: 233.252.0.1
MAC          MAC          Logical
address      flags          interface
00:00:5e:00:53:f4  D,SE          ge-4/2/0.1000
00:00:5e:00:53:33  D,SE          lsi.1052004
00:00:5e:00:53:32  D,SE          lsi.1048840
00:00:5e:00:53:14  D,SE          lsi.1052005
00:00:5e:00:53:f7  D,SE          vtep.1052010
00:00:5e:00:53:3f  D,SE          vtep.1052011
```

show vpls mac-table age (for GE interface)

```
user@host> show vpls mac-table age 00:00:5e:00:53:1a instance vpls_instance_1
```

```
MAC Entry Age information
Current Age: 4 seconds
```

show vpls mac-table age (for AE interface)

```
user@host> show vpls mac-table age 000:00:5e:00:53:1a instance vpls_instance_1
```

```
MAC Entry Age information
Current Age on FPC1: 102 seconds
Current Age on FPC2: 94 seconds
```

show vpls mac-table count

```
user@host> show vpls mac-table count
```

```
0 MAC address learned in routing instance __example_private1__
```

```
MAC address count per interface within routing instance:
```

Logical interface	MAC count
lc-0/0/0.32769	0
lc-0/1/0.32769	0
lc-0/2/0.32769	0
lc-2/0/0.32769	0
lc-0/3/0.32769	0
lc-2/1/0.32769	0
lc-9/0/0.32769	0
lc-11/0/0.32769	0
lc-2/2/0.32769	0
lc-9/1/0.32769	0
lc-11/1/0.32769	0
lc-2/3/0.32769	0
lc-9/2/0.32769	0
lc-11/2/0.32769	0
lc-11/3/0.32769	0
lc-9/3/0.32769	0

```
MAC address count per learn VLAN within routing instance:
```

Learn VLAN ID	MAC count
0	0

```
1 MAC address learned in routing instance vpls_ldp1
```

```
MAC address count per interface within routing instance:
```

Logical interface	MAC count
lsi.1051137	0
ge-0/2/5.400	1

```
MAC address count per learn VLAN within routing instance:
```

Learn VLAN ID	MAC count
0	1

```
1 MAC address learned in routing instance vpls_red
```

```
MAC address count per interface within routing instance:
```

Logical interface	MAC count
ge-0/2/5.300	1

```
MAC address count per learn VLAN within routing instance:
```

Learn VLAN ID	MAC count
0	1

show vpls mac-table detail

```
user@host> show vpls mac-table detail
```

```
MAC address: 00:00:5e:00:53:5d
Routing instance: vpls_ldpl
Learning interface: ge-0/2/5.400
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 0                               Sequence number: 1
Learning mask: 0x1                      IPC generation: 0
```

```
MAC address: 00:00:5e:00:53:5d
Routing instance: vpls_red
Learning interface: ge-0/2/5.300
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 0                               Sequence number: 1
Learning mask: 0x1                      IPC generation: 0
```

show vpls mac-table extensive

```
user@host> show vpls mac-table extensive
```

```
MAC address: 00:00:5e:00:53:00
Routing instance: vpls_1
Bridging domain: __vpls_1__, VLAN : NA
Learning interface: lsi.1049165
Base learning interface: lsi.1049165
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0                               Sequence number: 1
Learning mask: 0x00000001

MAC address: 00:00:5e:00:53:01
```

```

Routing instance: vpls_1
  Bridging domain: __vpls_1__, VLAN : NA
  Learning interface: lsi.1049165
  Base learning interface: lsi.1049165
  Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
  Epoch: 0                               Sequence number: 1
  Learning mask: 0x00000001

```

MAC address: 00:00:5e:00:53:02

```

Routing instance: vpls_1
  Bridging domain: __vpls_1__, VLAN : NA
  Learning interface: lsi.1049165
  Base learning interface: lsi.1049165
  Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
  Epoch: 0                               Sequence number: 1
  Learning mask: 0x00000001

```

MAC address: 00:00:5e:00:53:03

```

Routing instance: vpls_1
  Bridging domain: __vpls_1__, VLAN : NA
  Learning interface: lsi.1049165
  Base learning interface: lsi.1049165
  Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
  Epoch: 0                               Sequence number: 1
  Learning mask: 0x00000001

```

show vpls statistics

Syntax

```
show vpls statistics
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Release Information

Command introduced before Junos OS Release 7.4.

Description

(T Series and M Series routers, except for the M160 router) Display virtual private LAN service (VPLS) statistics.

Options

none—Display VPLS statistics for all routing instances.

instance *instance-name*—(Optional) Display VPLS statistics for a specific VPLS routing instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

[show vpls statistics on page 938](#)

[show vpls statistics instance on page 939](#)

Output Fields

[Table 48 on page 937](#) lists the output fields for the **show vpls statistics** command. Output fields are listed in the approximate order in which they appear.

Table 48: show vpls statistics Output Fields

Field Name	Field Description
Instance	Name of the VPLS instance.
Local interface	Name of the local VPLS virtual loopback tunnel interface, vt-fpc/pic/port.nnnnn , where nnnnn is a dynamically generated virtual port used to transport and receive packets from other provider edge (PE) routers in the VPLS domain.

Table 48: show vpls statistics Output Fields (continued)

Field Name	Field Description
Index	Number associated with the next hop.
Remote provider edge router	Address of the remote provider edge router.
Multicast packets	Number of multicast packets received.
Multicast bytes	Number of multicast bytes received.
Flood packets	Number of VPLS flood packets received.
Flood bytes	Number of VPLS flood bytes received.
Current MAC count	Number of MAC addresses learned by the interface and the configured maximum limit on the number of MAC addresses that can be learned.

Sample Output

show vpls statistics

user@host> **show vpls statistics**

VPLS statistics:

Instance: green

Local interface: fe-2/2/1.0, Index: 69

```

Multicast packets:           1
Multicast bytes   :          60
Flooded packets   :          18
Flooded bytes    :         2556
Current MAC count:           1

```

Local interface: lt-0/3/0.2, Index: 72

```

Multicast packets:           3
Multicast bytes   :         153
Flooded packets   :           1

```

```

        Flooded bytes      :                51
        Current MAC count:                1

Local interface: lsi.32769, Index: 75
        Current MAC count:                0

Local interface: lsi.32771, Index: 77
Remote PE: 10.255.14.222
        Current MAC count:                2

Instance: red

Local interface: vt-0/3/0.32768, Index: 74
        Multicast packets:                0
        Multicast bytes   :                0
        Flooded packets   :                0
        Flooded bytes     :                0
        Current MAC count:                0

Local interface: vt-0/3/0.32770, Index: 76
        Multicast packets:                0
        Multicast bytes   :                0
        Flooded packets   :                0
        Flooded bytes     :                0
        Current MAC count:                0

```

show vpls statistics instance

user@host> show vpls statistics instance red

```

Layer-2 VPN Statistics:
Instance: red

Local interface: vt-3/2/0.32768, Index: 73
Remote provider edge router: 10.255.17.35
        Multicast packets:                0
        Multicast bytes   :                0
        Flood packets     :                0
        Flood bytes       :                0
        Current MAC count:                1 (Limit 20)

```