



Junos[®] OS

Interfaces Feature Guide for Security Devices

Release

15.1X49-D70



Modified: 2016-11-15

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

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

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

Junos[®] OS Interfaces Feature Guide for Security Devices
15.1X49-D70
Copyright © 2016, Juniper Networks, Inc.
All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

	About the Documentation	xvii
	Documentation and Release Notes	xvii
	Supported Platforms	xvii
	Using the Examples in This Manual	xvii
	Merging a Full Example	xviii
	Merging a Snippet	xviii
	Documentation Conventions	xix
	Documentation Feedback	xxi
	Requesting Technical Support	xxi
	Self-Help Online Tools and Resources	xxi
	Opening a Case with JTAC	xxii
Part 1	Overview	
Chapter 1	Introduction to Interfaces	3
	Understanding Interfaces	3
	Network Interfaces	4
	Services Interfaces	5
	Special Interfaces	8
	Interface Naming Conventions	8
	Understanding the Data Link Layer	9
	Physical Addressing	10
	Network Topology	10
	Error Notification	10
	Frame Sequencing	10
	Flow Control	10
	Data Link Sublayers	11
	MAC Addressing	11
	Monitoring Interfaces	12
Chapter 2	Configuring Interface Logical Properties	15
	Understanding Interface Logical Properties	15
	Understanding Protocol Families	16
	Common Protocol Suites	16
	Other Protocol Suites	16
	Understanding IPv4 Addressing	17
	IPv4 Classful Addressing	17
	IPv4 Dotted Decimal Notation	18
	IPv4 Subnetting	18
	IPv4 Variable-Length Subnet Masks	19

	Understanding IPv6 Address Space, Addressing, Address Format, and Address Types	19
	Understanding IP Version 6 (IPv6)	20
	Understanding IPv6 Address Types and How Junos OS for SRX Series Services Gateway Uses Them	20
	IPv6 Address Scope	21
	IPv6 Address Structure	21
	Understanding IPv6 Address Space, Addressing, and Address Types	22
	Understanding IPv6 Address Format	22
	Limitations	23
	Configuring the inet6 IPv6 Protocol Family	23
	Enabling Flow-Based Processing for IPv6 Traffic	24
	Configuring Flow Aggregation to Use Version 9 Flow Templates	25
	Configuring the Traffic to Be Sampled	26
	Configuring the Version 9 Template Properties	26
	Restrictions	28
	Fields Included in Each Template Type	29
	inet Sampling Behavior	30
	Verification	31
	Examples: Configuring Version 9 Flow Templates	31
	Understanding IPv6 Support VDSL2 Interfaces	34
	Understanding MAC Limiting on Layer 3 Routing Interfaces	35
	Overview	35
	Limitations	37
Chapter 3	Understanding Interface Physical Properties	39
	Understanding Interface Physical Properties	39
	Understanding Bit Error Rate Testing	40
	Understanding Interface Clocking	41
	Data Stream Clocking	42
	Explicit Clocking Signal Transmission	42
	Understanding Frame Check Sequences	42
	Cyclic Redundancy Checks and Checksums	42
	Two-Dimensional Parity	43
	MTU Default and Maximum Values	43
Chapter 4	Configuring VLAN Tagging	47
	Understanding Virtual LANs	47
	VLAN IDs and Ethernet Interface Types Supported on the SRX Series Devices	49
	Configuring VLAN Tagging	49
	Configuring Single-Tag Framing	50
	Configuring Dual Tagging	50
	Configuring Mixed Tagging	50
	Configuring Mixed Tagging Support for Untagged Packets	51

Part 2	Configuring DS1 Interfaces	
Chapter 5	Configuring DS1 Interfaces	55
	Understanding T1 and E1 Interfaces	55
	T1 Overview	55
	E1 Overview	56
	T1 and E1 Signals	56
	Encoding	56
	AMI Encoding	56
	B8ZS and HDB3 Encoding	57
	T1 and E1 Framing	57
	ESF Framing for T1	57
	T1 and E1 Loopback Signals	57
	Example: Configuring a T1 Interface	58
	Example: Deleting a T1 Interface	60
Chapter 6	Configuring DS3 Interfaces	63
	Understanding T3 and E3 Interfaces	63
	Multiplexing DS1 Signals	63
	DS2 Bit Stuffing	64
	DS3 Framing	64
	M13 Asynchronous Framing	64
	C-Bit Parity Framing	66
	Example: Configuring a T3 Interface	68
	Example: Deleting a T3 Interface	70
Part 3	Configuring DSL Interfaces	
Chapter 7	Configuring VDSL2 Interfaces	75
	VDSL2 Interface Technology Overview	75
	VDSL2 Vectoring Overview	76
	VDSL2 Network Deployment Topology	76
	VDSL2 Interface Support on SRX Series Devices	77
	VDSL2 Interface Compatibility with ADSL Interfaces	78
	VDSL2 Interfaces Supported Profiles	79
	VDSL2 Interfaces Supported Features	80
	Example: Configuring VDSL2 Interfaces (Basic)	81
	Example: Configuring VDSL2 Interfaces (Detail)	87
	Upgrading the VDSL PIC Firmware	112
Part 4	Configuring Ethernet Interfaces	
Chapter 8	Performing Initial Configuration on Ethernet Interfaces	117
	Understanding Ethernet Interfaces	117
	Ethernet Access Control and Transmission	118
	Collisions and Detection	118
	Collision Detection	118
	Backoff Algorithm	118

	Collision Domains and LAN Segments	119
	Repeaters	119
	Bridges and Switches	119
	Broadcast Domains	120
	Ethernet Frames	120
	Understanding Static ARP Entries on Ethernet Interfaces	121
	Understanding Promiscuous Mode on Ethernet Interface	121
	Understanding Promiscuous Mode on the SRX5K-MPC	122
	Example: Creating an Ethernet Interface	122
	Example: Deleting an Ethernet Interface	123
	Example: Configuring Static ARP Entries on Ethernet Interfaces	124
	Enabling and Disabling Promiscuous Mode on Ethernet Interfaces (CLI Procedure)	127
	Example: Configuring Promiscuous Mode on the SRX5K-MPC	127
Chapter 9	Configuring Aggregated Ethernet Interfaces	133
	Understanding Aggregated Ethernet Interfaces	133
	LAGs	134
	LACP	134
	Aggregated Ethernet Interfaces Configuration Overview	136
	Understanding the Aggregated Ethernet Interfaces Device Count	136
	Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device	137
	Understanding Physical Interfaces for Aggregated Ethernet Interfaces	138
	Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces	138
	Understanding Aggregated Ethernet Interface Link Speed	139
	Example: Configuring Aggregated Ethernet Link Speed	140
	Understanding Minimum Links for Aggregated Ethernet Interfaces	141
	Example: Configuring Aggregated Ethernet Minimum Links	141
	Understanding Aggregated Ethernet Interface Removal	142
	Example: Deleting Aggregated Ethernet Interfaces	143
	Example: Deleting Aggregated Ethernet Interface Contents	144
	Verifying Aggregated Ethernet Interfaces	145
	Verifying Aggregated Ethernet Interfaces (terse)	145
	Verifying Aggregated Ethernet Interfaces (extensive)	145
	Understanding VLAN Tagging for Aggregated Ethernet Interfaces	146
	Understanding Promiscuous Mode for Aggregated Ethernet Interfaces	147
Chapter 10	Configuring Link Aggregation Control Protocol	149
	Understanding LACP on Standalone Devices	149
	Example: Configuring LACP on Standalone Devices	150
	Verifying LACP on Standalone Devices	151
	Verifying LACP Statistics	151
	Verifying LACP Aggregated Ethernet Interfaces	152
	Understanding LACP on Chassis Clusters	153
	Chassis Cluster Redundant Ethernet Interface Link Aggregation Groups	153
	Minimum Links	154
	Sub-LAGs	154
	Supporting Hitless Failover	154

	Managing Link Aggregation Control PDUs	155
	Example: Configuring LACP on Chassis Clusters	155
	Verifying LACP on Redundant Ethernet Interfaces	157
	LAG and LACP Support on SRX5000 Line Devices with I/O Cards (IOCs)	158
	LAG and LACP Support on the SRX5000 Module Port Concentrator	158
	LAG and LACP Support on the SRX5000 Line IOCs in Express Path Mode	159
	Example: Configuring LAG Interface on an SRX5000 Line Device with IOC2 or IOC3	160
Chapter 11	Configuring Gigabit Ethernet Physical Interface Modules	165
	Understanding the 1-Port Gigabit Ethernet SFP Mini-PIM	165
	Supported Features	165
	Interface Names and Settings	166
	Available Link Speeds and Modes	166
	Link Settings	167
	Example: Configuring the 1-Port Gigabit Ethernet SFP Mini-PIM Interface	167
	Understanding the 2-Port 10-Gigabit Ethernet XPIM	173
	Supported Features	174
	Interface Names and Settings	174
	Copper and Fiber Operating Modes	175
	Link Speeds	175
	Link Settings	175
	Example: Configuring the 2-Port 10-Gigabit Ethernet XPIM Interface	176
Chapter 12	Configuring Ethernet OAM Link Fault Management	181
	Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways	181
	Example: Configuring Ethernet OAM Link Fault Management	183
Chapter 13	Configuring Power over Ethernet	187
	Understanding Power over Ethernet	187
	SRX Series Services Gateway PoE Specifications	187
	PoE Classes and Power Ratings	188
	PoE Options	189
	Example: Configuring PoE on All Interfaces	189
	Example: Configuring PoE on an Individual Interface	191
	Example: Disabling a PoE Interface	194
Part 5	Configuring Interface Encapsulation	
Chapter 14	Interface Encapsulation Overview	199
	Understanding Physical Encapsulation on an Interface	199
	Understanding Frame Relay Encapsulation on an Interface	200
	Virtual Circuits	200
	Switched and Permanent Virtual Circuits	201
	Data-Link Connection Identifiers	201
	Congestion Control and Discard Eligibility	201

	Understanding Point-to-Point Protocol	202
	Link Control Protocol	202
	PPP Authentication	203
	Network Control Protocols	203
	Magic Numbers	204
	CSU/DSU Devices	204
	Understanding High-Level Data Link Control	204
	HDLC Stations	205
	HDLC Operational Modes	205
Chapter 15	Configuring Point-to-Point Protocol over Ethernet	207
	Understanding Point-to-Point Protocol over Ethernet	207
	PPPoE Discovery Stage	208
	PPPoE Session Stage	209
	Understanding PPPoE Interfaces	210
	Example: Configuring PPPoE Interfaces	211
	Understanding PPPoE Ethernet Interfaces	217
	Example: Configuring PPPoE Encapsulation on an Ethernet Interface	217
	Understanding PPPoE ATM-over-ADSL and ATM-over-SHDSL Interfaces	218
	Example: Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface	219
	Understanding CHAP Authentication on a PPPoE Interface	221
	Example: Configuring CHAP Authentication on a PPPoE Interface	221
	Verifying Credit-Flow Control	223
	Verifying PPPoE Interfaces	224
	Verifying R2CP Interfaces	225
	Displaying Statistics for PPPoE	226
	Setting Tracing Options for PPPoE	226
Chapter 16	Configuring PPPoE-Based Radio-to-Router Protocol	229
	PPPoE-Based Radio-to-Router Protocols Overview	229
	Understanding the PPPoE-Based Radio-to-Router Protocol	230
	Configuring PPPoE-Based Radio-to-Router Protocols	232
	Example: Configuring the PPPoE-Based Radio-to-Router Protocol	232
	Credit Flow Control for PPPoE	235
	PPPoE Credit-Based Flow Control Configuration	235
Chapter 17	Configuring R2CP Radio-to-Router Protocol	237
	R2CP Radio-to-Router Protocol Overview	237
	Configuring the R2CP Radio-to-Router Protocol	238
Part 6	Configuration Statements and Operational Commands	
Chapter 18	Configuration Statements	245
	accept-source-mac	248
	access-point-name	249
	apply-groups	249
	arp-resp	250
	authentication-method (Interfaces)	250
	bandwidth (Interfaces)	251
	bundle (Interfaces)	251

cbr rate	252
cellular-options	252
classifiers (CoS)	253
client-identifier (Interfaces)	254
code-points (CoS)	254
compression-device (Interfaces)	255
credit (Interfaces)	255
data-rate	256
disable (PoE)	256
dhcp (Interfaces)	257
duration (PoE)	258
encapsulation (Interfaces)	259
family inet (Interfaces)	260
family inet6	263
flag (Interfaces)	265
flexible-vlan-tagging (Interfaces)	266
flow-control (Interfaces)	266
flow-monitoring (Services)	267
forwarding-classes (CoS)	268
fpc (Interfaces)	269
gratuitous-arp-reply	270
gsm-options	270
guard-band (PoE)	271
hub-assist	271
inline-jflow (Forwarding Options)	272
interface (PIC Bundle)	272
interface (PoE)	273
interfaces (CoS)	274
interval (Interfaces)	275
interval (PoE)	275
ipv4-template (Services)	276
ipv6-template (Services)	276
lACP (Interfaces)	277
latency (Interfaces)	277
lease-time	278
line-rate (Interfaces)	278
link-speed (Interfaces)	279
loopback (Interfaces)	279
loss-priority (CoS Loss Priority)	280
loss-priority (CoS Rewrite Rules)	281
loss-priority-maps (CoS Interfaces)	282
loss-priority-maps (CoS)	282
management (PoE)	283
maximum-power (PoE)	283
media-type (Interfaces)	284
minimum-links (Interfaces)	285
native-vlan-id (Interfaces)	286
next-hop-tunnel	286
no-dns-propagation	287

option-refresh-rate (Services)	287
pic-mode (Chassis T1 Mode)	288
periodic (Interfaces)	289
ppp-over-ether	289
pppoe	290
pppoe-options	291
priority (PoE)	292
profile (Access)	293
profiles	295
promiscuous-mode (Interfaces)	296
quality (Interfaces)	296
r2cp	297
radio-router (Interfaces)	298
redundancy-group (Interfaces)	299
redundant-ether-options	300
redundant-parent (Interfaces Fast Ethernet)	301
redundant-parent (Interfaces Gigabit Ethernet)	301
resource (Interfaces)	302
retransmission-attempt	302
retransmission-interval (Interfaces)	303
roaming-mode	303
scheduler-map (CoS Virtual Channels)	304
select-profile	304
server-address	305
shaping-rate (CoS Interfaces)	306
simple-filter (Interfaces)	307
sip-password	307
sip-user-id	308
source-address-filter (Interfaces)	309
source-filtering (Interfaces)	310
speed (Interfaces)	310
telemetries (PoE)	311
template-refresh-rate (Services)	311
threshold (Interfaces)	312
traceoptions (Interfaces)	312
update-server	313
vbr rate	313
vdsl-profile	314
vendor-id (Interfaces)	314
vlan-tagging (Interfaces)	315
web-authentication (Interfaces)	316
Chapter 19	
Operational Commands	317
clear oam ethernet connectivity-fault-management path-database	319
clear dhcpv6 server binding (Local Server)	320
clear ethernet-switching statistics mac-learning	321
clear interfaces statistics swfabx	322
clear ipv6 neighbors	323
clear lacp statistics interfaces	324

restart (Reset)	325
show chassis fpc (View)	330
show chassis hardware (View)	338
show ethernet-switching mac-learning-log (View)	349
show ethernet-switching table (View)	351
show igmp-snooping route (View)	356
show interfaces (SRX Series)	358
show interfaces diagnostics optics	389
show interfaces flow-statistics	394
show interfaces queue	399
show interfaces statistics (View)	402
show interfaces terse zone	403
show ipv6 neighbors	404
show lacp interfaces (View)	406
show lacp statistics interfaces (View)	410
show oam ethernet link-fault-management	411
show poe controller (View)	416
show pppoe interfaces	417
show pppoe statistics	421
show poe telemetries	423
show services accounting	425
show services accounting aggregation (View)	428
show services accounting aggregation template (View)	429
show services accounting flow-detail (View)	430

List of Figures

Part 1	Overview	
Chapter 2	Configuring Interface Logical Properties	15
	Figure 1: Subnets in a Network	18
Chapter 4	Configuring VLAN Tagging	47
	Figure 2: Typical LAN	48
	Figure 3: Typical VLAN	48
Part 2	Configuring DS1 Interfaces	
Chapter 6	Configuring DS3 Interfaces	63
	Figure 4: DS2 M-Frame Format	64
	Figure 5: DS3 M13 Frame Format	65
	Figure 6: DS3 C-Bit Parity Framing	66
Part 3	Configuring DSL Interfaces	
Chapter 7	Configuring VDSL2 Interfaces	75
	Figure 7: Typical VDSL2 End-to-End Connectivity and Topology Diagram	77
	Figure 8: Backward-Compatible ADSL Topology (ATM DSLAM)	77
	Figure 9: SRX Series Device with VDSL2 Mini-PIMs in an End-to-End Deployment Scenario	88
Part 4	Configuring Ethernet Interfaces	
Chapter 8	Performing Initial Configuration on Ethernet Interfaces	117
	Figure 10: Ethernet Frame Format	120
Chapter 12	Configuring Ethernet OAM Link Fault Management	181
	Figure 11: Ethernet LFM with SRX Series Devices	184
Part 5	Configuring Interface Encapsulation	
Chapter 14	Interface Encapsulation Overview	199
	Figure 12: Frame Relay Network	200
Chapter 15	Configuring Point-to-Point Protocol over Ethernet	207
	Figure 13: PPPoE Session on the Ethernet Loop	217
	Figure 14: PPPoE Session on an ADSL Loop	219

List of Tables

	About the Documentation	xvii
	Table 1: Notice Icons	xix
	Table 2: Text and Syntax Conventions	xix
Part 1	Overview	
Chapter 1	Introduction to Interfaces	3
	Table 3: Network Interfaces	4
	Table 4: Configurable Services Interfaces	5
	Table 5: Non-Configurable Services Interfaces	7
	Table 6: Special Interfaces	8
	Table 7: Network Interface Names	9
Chapter 2	Configuring Interface Logical Properties	15
	Table 8: Device Status Upon Configuration Change	25
Chapter 3	Understanding Interface Physical Properties	39
	Table 9: Interface Physical Properties	39
	Table 10: MTU Values for the SRX Series Services Gateways PIMs	44
Chapter 4	Configuring VLAN Tagging	47
	Table 11: VLAN ID Range by Interface Type Supported on the SRX Series Devices	49
	Table 12: Flexible VLANs	49
Part 2	Configuring DS1 Interfaces	
Chapter 6	Configuring DS3 Interfaces	63
	Table 13: FEAC C-Bit Condition Indicators	67
Part 3	Configuring DSL Interfaces	
Chapter 7	Configuring VDSL2 Interfaces	75
	Table 14: VDSL2 Annex A and Annex B Features	78
	Table 15: VDSL2 Operating Mode Backward Compatibility with ADSL	79
	Table 16: Supported Profiles on the VDSL2 Interfaces	80
Part 4	Configuring Ethernet Interfaces	
Chapter 8	Performing Initial Configuration on Ethernet Interfaces	117
	Table 17: Collision Backoff Algorithm Rounds	118
Chapter 12	Configuring Ethernet OAM Link Fault Management	181

	Table 18: Supported Interface Modes	182
Chapter 13	Configuring Power over Ethernet	187
	Table 19: PoE Specifications for the SRX320 and SRX340 Devices	187
	Table 20: SRX Series Devices PoE Specifications	188
Part 6	Configuration Statements and Operational Commands	
Chapter 19	Operational Commands	317
	Table 21: SRX5K-MPC3-40G10G (IOC3) PIC Selection Summary	330
	Table 22: show chassis fpc Output Fields	332
	Table 23: show chassis hardware Output Fields	338
	Table 24: show interfaces Output Fields	349
	Table 25: show ethernet-switching table Output Fields	351
	Table 26: show igmp-snooping route Output Fields	356
	Table 27: show interfaces Output Fields	361
	Table 28: show interfaces diagnostics optics Output Fields	389
	Table 29: show interfaces flow-statistics Output Fields	395
	Table 30: Flow Error Statistics (Packet Drop Statistics for the Flow Module)	395
	Table 31: show interfaces queue Output Fields	399
	Table 32: show ipv6 neighbors Output Fields	404
	Table 33: show lacp interfaces Output Fields	406
	Table 34: show lacp statistics interfaces Output Fields	410
	Table 35: show oam ethernet link-fault-management Output Fields	411
	Table 36: show poe controller Output Fields	416
	Table 37: show pppoe interfaces Output Fields	417
	Table 38: show pppoe statistics Output Fields	421
	Table 39: show poe telemetries interface Output Fields	423

About the Documentation

- Documentation and Release Notes on page xvii
- Supported Platforms on page xvii
- Using the Examples in This Manual on page xvii
- Documentation Conventions on page xix
- Documentation Feedback on page xxi
- Requesting Technical Support on page xxi

Documentation and Release Notes

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

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

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

Supported Platforms

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

- SRX Series
- vSRX

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

- 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 [xix](#) 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 [xix](#) 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

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

Convention	Description	Examples
Fixed-width text like this	Represents output that appears on the terminal screen.	<code>user@host> show chassis alarms</code> <code>No alarms currently active</code>
<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>
<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: <code>[edit]</code> <code>root@# set system domain-name <i>domain-name</i></code>
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 <code>stub</code> statement at the <code>[edit protocols ospf area area-id]</code> hierarchy level. The console port is labeled <code>CONSOLE</code>.
< > (angle brackets)	Encloses optional keywords or variables.	<code>stub <default-metric <i>metric</i>>;</code>
(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.	<code>broadcast multicast</code> <code>(<i>string1</i> <i>string2</i> <i>string3</i>)</code>
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	<code>rsvp { # Required for dynamic MPLS only</code>
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	<code>community name members [<i>community-ids</i>]</code>
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	<code>[edit]</code> <code>routing-options {</code> <code> static {</code> <code> route default {</code> <code> nexthop <i>address</i>;</code> <code> retain;</code> <code> }</code> <code> }</code> <code>}</code>
:(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> In the Logical Interfaces box, select All Interfaces. To cancel the configuration, click Cancel.

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

Convention	Description	Examples
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

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

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

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>

- Download the latest versions of software and review release notes:
<http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications:
<http://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum:
<http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

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

Opening a Case with JTAC

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

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

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

PART 1

Overview

- [Introduction to Interfaces on page 3](#)
- [Configuring Interface Logical Properties on page 15](#)
- [Understanding Interface Physical Properties on page 39](#)
- [Configuring VLAN Tagging on page 47](#)

CHAPTER 1

Introduction to Interfaces

- [Understanding Interfaces on page 3](#)
- [Network Interfaces on page 4](#)
- [Services Interfaces on page 5](#)
- [Special Interfaces on page 8](#)
- [Interface Naming Conventions on page 8](#)
- [Understanding the Data Link Layer on page 9](#)
- [Monitoring Interfaces on page 12](#)

Understanding Interfaces

Supported Platforms [SRX Series, vSRX](#)

Interfaces act as a doorway through which traffic enters and exits a device. Juniper Networks devices support a variety of interface types:

- Network interfaces—Networking interfaces primarily provide traffic connectivity.
- Services interfaces—Services interfaces manipulate traffic before it is delivered to its destination.
- Special interfaces—Special interfaces include management interfaces, the loopback interface, and the discard interface.

Each type of interface uses a particular medium to transmit data. The physical wires and Data Link Layer protocols used by a medium determine how traffic is sent. To configure and monitor interfaces, you need to understand their media characteristics, as well as physical and logical properties such as IP addressing, link-layer protocols, and link encapsulation.



NOTE: Most interfaces are configurable, but some internally generated interfaces are not configurable.

Related Documentation

- [Interface Naming Conventions on page 8](#)
- [Understanding Interface Logical Properties on page 15](#)

- [Understanding Interface Physical Properties](#) on page 39
- [Understanding the Data Link Layer](#) on page 9

Network Interfaces

Supported Platforms [SRX Series, vSRX](#)

All Juniper Networks devices use network interfaces to make physical connections to other devices. A connection takes place along media-specific physical wires through an I/O card (IOC) in the SRX Series Services Gateway. Networking interfaces primarily provide traffic connectivity.

You must configure each network interface before it can operate on the device. Configuring an interface can define both the physical properties of the link and the logical properties of a logical interface on the link.

[Table 3 on page 4](#) describes network interfaces that are available on SRX Series devices.

Table 3: Network Interfaces

Interface Name	Description
ae	Aggregated Ethernet interface. See "Understanding Aggregated Ethernet Interfaces" on page 133.
at	ATM-over-ADSL or ATM-over-SHDSL WAN interface.
dl	Dialer interface for initiating USB modem connections. See USB Modem Interface Overview .
e1	E1 (also called DS1) WAN interface. See "Understanding T1 and E1 Interfaces" on page 55.
e3	E3 (also called DS3) WAN interface. See "Understanding T3 and E3 Interfaces" on page 63.
fe	Fast Ethernet interface. See "Understanding Ethernet Interfaces" on page 117.
ge	Gigabit Ethernet interface. See "Understanding Ethernet Interfaces" on page 117.
pt	VDSL2 interface. See "Example: Configuring VDSL2 Interfaces (Detail)" on page 87.
reth	For chassis cluster configurations only, redundant Ethernet interface. See "Understanding Ethernet Interfaces" on page 117.
se	Serial interface (either RS-232, RS-422/499, RS-530, V.35, or X.21).
t1	T1 (also called DS1) WAN interface. See "Understanding T1 and E1 Interfaces" on page 55.
t3	T3 (also called DS3) WAN interface. See "Understanding T3 and E3 Interfaces" on page 63.
wx	WXC Integrated Services Module (ISM 200) interface for WAN acceleration. See the WXC Integrated Services Module Installation and Configuration .

Table 3: Network Interfaces (*continued*)

Interface Name	Description
xe	10-Gigabit Ethernet interface. See "Understanding the 2-Port 10-Gigabit Ethernet XPIM" on page 173.

- Related Documentation**
- [Understanding Interfaces on page 3](#)
 - [Services Interfaces on page 5](#)
 - [Special Interfaces on page 8](#)

Services Interfaces

Supported Platforms [SRX Series, vSRX](#)

Services interfaces provide specific capabilities for manipulating traffic before it is delivered to its destination. On Juniper Networks M Series and T Series routing platforms, individual services such as IP-over-IP encapsulation, link services such as multilink protocols, adaptive services such as stateful firewall filters and NAT, and sampling and logging capabilities are implemented by services Physical Interface Cards (PICs). On SRX Series devices, services processing is handled by the Services Processing Card (SPC).

Although the same Junos OS image supports the services features across all routing platforms, on SRX Series devices, services interfaces are not associated with a physical interface. To configure services on these devices, you configure one or more internal interfaces by specifying slot **0**, interface carrier **0**, and port **0**—for example, **gr-0/0/0** for GRE.

[Table 4 on page 5](#) describes services interfaces that you can configure on SRX Series Services Gateways.

Table 4: Configurable Services Interfaces

Interface Name	Description
gr-0/0/0	<p>Configurable generic routing encapsulation (GRE) interface. GRE allows the encapsulation of one routing protocol inside another routing protocol.</p> <p>Packets are routed to this internal interface, where they are first encapsulated with a GRE packet and then sent.</p> <p>You can create multiple instances of this interface for forwarding encapsulated data to multiple destination addresses by using the default interface as the parent and creating extensions, for example, gr-0/0/0.1, gr-0/0/0.2, and so on.</p> <p>The GRE interface is an internal interface only and is not associated with a physical interface. It is used only for processing GRE traffic. See the Junos OS Services Interfaces Library for Routing Devices for information about tunnel services.</p>

Table 4: Configurable Services Interfaces (*continued*)

Interface Name	Description
ip-0/0/0	<p>Configurable IP-over-IP encapsulation (IP-IP tunnel) interface. IP tunneling allows the encapsulation of one IP packet inside another IP packet.</p> <p>With IP routing, you can route IP packets directly to a particular address or route the IP packets to an internal interface where they are encapsulated inside an IP-IP tunnel and forwarded to the encapsulating packet's destination address.</p> <p>You can create multiple instances of this interface for forwarding IP-IP tunnel data to multiple destination addresses by using the default interface as the parent and creating extensions, for example, ip-0/0/0.1, ip-0/0/0.2, and so on.</p> <p>The IP-IP interface is an internal interface only and is not associated with a physical interface. It is used only for processing IP-IP tunnel traffic. See the Junos OS Services Interfaces Library for Routing Devices for information about tunnel services.</p>
lsq-0/0/0	<p>Configurable link services queuing interface. Link services include the multilink services MLPPP, MLFR, and Compressed Real-Time Transport Protocol (CRTP).</p> <p>Packets are routed to this internal interface for link bundling or compression. The link services interface is an internal interface only and is not associated with a physical interface. You must configure the interface for it to perform multilink services.</p> <p>NOTE: The ls-0/0/0 interface has been deprecated. All multiclass multilink features supported by ls-0/0/0 are now supported by lsq-0/0/0.</p>
lt-0/0/0	<p>Configurable logical tunnel interface that interconnects logical systems on SRX Series devices. See the <i>Logical Systems Feature Guide for Security Devices</i>.</p>
pp0	<p>Configurable PPPoE encapsulation interface. PPP packets being routed in an Ethernet network use PPPoE encapsulation.</p> <p>Packets are routed to this internal interface for PPPoE encapsulation. The PPPoE encapsulation interface is an internal interface only and is not associated with a physical interface. You must configure the interface for it to forward PPPoE traffic.</p> <p>See “Understanding Point-to-Point Protocol over Ethernet” on page 207.</p>
ppd0	<p>Protocol Independent Multicast (PIM) de-encapsulation interface. In PIM sparse mode, the first-hop routing platform encapsulates packets destined for the rendezvous point device. The packets are encapsulated with a unicast header and are forwarded through a unicast tunnel to the rendezvous point. The rendezvous point then de-encapsulates the packets and transmits them through its multicast tree.</p> <p>Within a device, packets are routed to this internal interface for de-encapsulation. The PIM de-encapsulation interface is an internal interface only and is not associated with a physical interface. You must configure PIM with the <code>[edit protocol pim]</code> hierarchy to perform PIM de-encapsulation.</p> <p>Use the <code>show pim interfaces</code> command to check the status of ppd0 interface.</p>

Table 4: Configurable Services Interfaces (*continued*)

Interface Name	Description
ppe0	<p>Protocol Independent Multicast (PIM) encapsulation interface. In PIM sparse mode, the first-hop routing platform encapsulates packets destined for the rendezvous point device. The packets are encapsulated with a unicast header and are forwarded through a unicast tunnel to the rendezvous point. The rendezvous point then de-encapsulates the packets and transmits them through its multicast tree.</p> <p>Within a device, packets are routed to this internal interface for encapsulation. The PIM encapsulation interface is an internal interface only and is not associated with a physical interface. You must configure PIM with the [edit protocol pim] hierarchy to perform PIM encapsulation.</p>
st0	Secure tunnel interface used for IPSec VPNs. See the <i>VPN Feature Guide for Security Devices</i> .
umd0	<p>Configurable USB modem physical interface. This interface is detected when a USB modem is connected to the USB port on the device.</p> <p>See <i>USB Modem Configuration Overview</i>.</p>

[Table 5 on page 7](#) describes non-configurable services interfaces for SRX Series Services Gateways.

Table 5: Non-Configurable Services Interfaces

Interface Name	Description
gre	Internally generated Generic Routing Encapsulation (GRE) interface created by Junos OS to handle GRE traffic. It is not a configurable interface.
ipip	Internally generated IP-over-IP interface created by Junos OS to handle IP tunnel traffic. It is not a configurable interface.
lsi	Internally generated link services interface created by Junos OS to handle multilink services like MLPPP, MLFR, and CRTP. It is not a configurable interface.
pc-pim/0/0	Internally configured interface used by the system as a control path between the WXC Integrated Services Module and the Routing Engine. It is not a configurable interface. See the WX and WXC Series .
pimd	Internally generated Protocol Independent Multicast (PIM) de-encapsulation interface created by Junos OS to handle PIM de-encapsulation. It is not a configurable interface.
pime	Internally generated Protocol Independent Multicast (PIM) encapsulation interface created by Junos OS to handle PIM encapsulation. It is not a configurable interface.
tap	Internally generated interface created by Junos OS to monitor and record traffic during passive monitoring. Packets discarded by the Packet Forwarding Engine are placed on this interface. It is not a configurable interface.

- Related Documentation**
- [Junos Services Interfaces Configuration](#)
 - [Understanding Interfaces on page 3](#)

- [Network Interfaces on page 4](#)
- [Special Interfaces on page 8](#)

Special Interfaces

Supported Platforms [SRX Series, vSRX](#)

Special interfaces include management interfaces, which are primarily intended for accessing the device remotely, the loopback interface, which has several uses depending on the particular Junos OS feature being configured, and the discard interface.

[Table 6 on page 8](#) describes special interfaces for SRX Series Services Gateways.

Table 6: Special Interfaces

Interface Name	Description
fxp0, fxp1	On SRX Series devices, the fxp0 management interface is a dedicated port located on the Routing Engine.
lo0	Loopback address. The loopback address has several uses, depending on the particular Junos feature being configured.
dsc	Discard interface.

- Related Documentation**
- [Understanding Interfaces on page 3](#)
 - [Network Interfaces on page 4](#)
 - [Services Interfaces on page 5](#)

Interface Naming Conventions

Supported Platforms [SRX Series, vSRX](#)

Each device interface has a unique name that follows a naming convention. If you are familiar with Juniper Networks M Series and T Series routing platforms, be aware that device interface names are similar to but not identical to the interface names on those routing platforms.

The unique name of each network interface identifies its type and location and indicates whether it is a physical interface or an optional logical unit created on a physical interface.

- The name of each network interface has the following format to identify the physical device that corresponds to a single physical network connector:

type-slot/pim-or-ioc/port

- Network interfaces that are fractionalized into time slots include a channel number in the name, preceded by a colon (:):

type-slot/pim-or-ioc/port:channel

- Each logical interface has an additional logical unit identifier, preceded by a period (.):

type-slot/pim-or-ioc/port:<channel>.unit

The parts of an interface name are summarized in [Table 7 on page 9](#).

Table 7: Network Interface Names

Name Part	Meaning	Possible Values
<i>type</i>	Type of network medium that can connect to this interface.	ae, at, ei, e3, fe, fxp0, fxp1, ge, lo0, lsq, lt, ppo, pt, sto, t1, t3, xe, and so on.
<i>slot</i>	Number of the chassis slot in which a PIM or IOC is installed.	SRX5600 and SRX5800 devices: The slot number begins at 0 and increases as follows from left to right, bottom to top: <ul style="list-style-type: none"> SRX5600 device—Slots 0 to 5 SRX5800 device—Slots 0 to 5, 7 to 11
<i>pim-or-ioc</i>	Number of the PIM or IOC on which the physical interface is located.	SRX5600 and SRX5800 devices: For 40-port Gigabit Ethernet IOCs or 4-port 10-Gigabit Ethernet IOCs, this number can be 0, 1, 2, or 3 .
<i>port</i>	Number of the port on a PIM or IOC on which the physical interface is located.	On SRX5600 and SRX5800 devices: <ul style="list-style-type: none"> For 40-port Gigabit Ethernet IOCs, this number begins at 0 and increases from left to right to a maximum of 9. For 4-port 10-Gigabit Ethernet IOCs, this number is always 0. Port numbers appear on the PIM or IOC faceplate.
<i>channel</i>	Number of the channel (time slot) on a fractional or channelized T1 or E1 interface.	<ul style="list-style-type: none"> On an E1 interface, a value from 1 through 31. The 1 time slot is reserved. On a T1 interface, a value from 1 through 24.
<i>unit</i>	Number of the logical interface created on a physical interface.	<p>A value from 0 through 16384.</p> <p>If no logical interface number is specified, unit 0 is the default, but must be explicitly configured.</p> <p>In addition to user-configured interfaces, there are some logical interfaces that are created dynamically. Hence, for Junos OS, the maximum limit for configuring logical interfaces is 2,62,143 (user configured and dynamically created). Based on performance, for each platform, the maximum number of logical interfaces supported can vary.</p>

Related Documentation • [Understanding Interfaces on page 3](#)

Understanding the Data Link Layer

Supported Platforms [SRX Series, vSRX](#)

The Data Link Layer is Layer 2 in the Open Systems Interconnection (OSI) model. The Data Link Layer is responsible for transmitting data across a physical network link. Each

physical medium has link-layer specifications for network and link-layer protocol characteristics such as physical addressing, network topology, error notification, frame sequencing, and flow control.

- [Physical Addressing on page 10](#)
- [Network Topology on page 10](#)
- [Error Notification on page 10](#)
- [Frame Sequencing on page 10](#)
- [Flow Control on page 10](#)
- [Data Link Sublayers on page 11](#)
- [MAC Addressing on page 11](#)

Physical Addressing

Physical addressing is different from network addressing. Network addresses differentiate between nodes or devices in a network, allowing traffic to be routed or switched through the network. In contrast, physical addressing identifies devices at the link-layer level, differentiating between individual devices on the same physical medium. The primary form of physical addressing is the media access control (MAC) address.

Network Topology

Network topology specifications identify how devices are linked in a network. Some media allow devices to be connected by a bus topology, while others require a ring topology. The bus topology is used by Ethernet technologies, which are supported on Juniper Networks devices.

Error Notification

The Data Link Layer provides error notifications that alert higher layer protocols that an error has occurred on the physical link. Examples of link-level errors include the loss of a signal, the loss of a clocking signal across serial connections, or the loss of the remote endpoint on a T1 or T3 link.

Frame Sequencing

The frame sequencing capabilities of the Data Link Layer allow frames that are transmitted out of sequence to be reordered on the receiving end of a transmission. The integrity of the packet can then be verified by means of the bits in the Layer 2 header, which is transmitted along with the data payload.

Flow Control

Flow control within the Data Link Layer allows receiving devices on a link to detect congestion and notify their upstream and downstream neighbors. The neighbor devices relay the congestion information to their higher layer protocols so that the flow of traffic can be altered or rerouted.

Data Link Sublayers

The Data Link Layer is divided into two sublayers: logical link control (LLC) and media access control (MAC). The LLC sublayer manages communications between devices over a single link of a network. This sublayer supports fields in link-layer frames that enable multiple higher layer protocols to share a single physical link.

The MAC sublayer governs protocol access to the physical network medium. Through the MAC addresses that are typically assigned to all ports on a device, multiple devices on the same physical link can uniquely identify one another at the Data Link Layer. MAC addresses are used in addition to the network addresses that are typically configured manually on ports within a network.

MAC Addressing

A MAC address is the serial number permanently stored in a device adapter to uniquely identify the device. MAC addresses operate at the Data Link Layer, while IP addresses operate at the Network Layer. The IP address of a device can change as the device is moved around a network to different IP subnets, but the MAC address remains the same, because it is physically tied to the device.

Within an IP network, devices match each MAC address to its corresponding configured IP address by means of the Address Resolution Protocol (ARP). ARP maintains a table with a mapping for each MAC address in the network.

Most Layer 2 networks use one of three primary numbering spaces—MAC-48, EUI-48 (extended unique identifier), and EUI-64—which are all globally unique. MAC-48 and EUI-48 spaces each use 48-bit addresses, and EUI-64 spaces use a 64-bit addresses, but all three use the same numbering format. MAC-48 addresses identify network hardware, and EUI-48 addresses identify other devices and software.

The Ethernet and ATM technologies supported on devices use the MAC-48 address space. IPv6 uses the EUI-64 address space.

MAC-48 addresses are the most commonly used MAC addresses in most networks. These addresses are 12-digit hexadecimal numbers (48 bits in length) that typically appear in one of the following formats:

- ***MM:MM:MM:SS:SS:SS***
- ***MM-MM-MM-SS-SS-SS***

The first three octets (***MM:MM:MM*** or ***MM-MM-MM***) are the ID number of the hardware manufacturer. Manufacturer ID numbers are assigned by the Institute of Electrical and Electronics Engineers (IEEE). The last three octets (***SS:SS:SS*** or ***SS-SS-SS***) make up the serial number for the device, which is assigned by the manufacturer. For example, an Ethernet interface card might have a MAC address of **00:05:85:c1:a6:a0**.

Related Documentation

- [Understanding Interfaces on page 3](#)

Monitoring Interfaces

Supported Platforms [SRX Series, vSRX](#)

Purpose View general information about all physical and logical interfaces for a device.

Action Select **Monitor > Interfaces** in the J-Web user interface. The J-Web Interfaces page displays the following details about each device interface:

- Port—Indicates the interface name.
- Admin Status—Indicates whether the interface is enabled (Up) or disabled (Down).
- Link Status—Indicates whether the interface is linked (Up) or not linked (Down).
- Address—Indicates the IP address of the interface.
- Zone—Indicates whether the zone is an untrust zone or a trust zone.
- Services—Indicates services that are enabled on the device, such as HTTP and SSH.
- Protocols—Indicates protocols that are enabled on the device, such as BGP and IGMP.
- Input Rate graph—Displays interface bandwidth utilization. Input rates are shown in bytes per second.
- Output Rate graph—Displays interface bandwidth utilization. Output rates are shown in bytes per second.
- Error Counters chart—Displays input and output error counters in the form of a bar chart.
- Packet Counters chart—Displays the number of broadcast, unicast, and multicast packet counters in the form of a pie chart. (Packet counter charts are supported only for interfaces that support MAC statistics.)

To change the interface display, use the following options:

- Port for FPC—Controls the member for which information is displayed.
- Start/Stop button—Starts or stops monitoring the selected interfaces.
- Show Graph—Displays input and output packet counters and error counters in the form of charts.
- Pop-up button—Displays the interface graphs in a separate pop-up window.
- Details—Displays extensive statistics about the selected interface, including its general status, traffic information, IP address, I/O errors, class-of-service data, and statistics.
- Refresh Interval—Indicates the duration of time after which you want the data on the page to be refreshed.
- Clear Statistics—Clears the statistics for the selected interface.

Alternatively, you can enter the following **show** commands in the CLI to view interface status and traffic statistics:

- **show interfaces terse**



NOTE: On SRX Series devices, on configuring identical IPs on a single interface, you will not see a warning message; instead, you will see a syslog message.

- **show interfaces detail**
- **show interfaces extensive**
- **show interfaces *interface-name***

CHAPTER 2

Configuring Interface Logical Properties

- [Understanding Interface Logical Properties on page 15](#)
- [Understanding Protocol Families on page 16](#)
- [Understanding IPv4 Addressing on page 17](#)
- [Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 19](#)
- [Configuring the inet6 IPv6 Protocol Family on page 23](#)
- [Enabling Flow-Based Processing for IPv6 Traffic on page 24](#)
- [Configuring Flow Aggregation to Use Version 9 Flow Templates on page 25](#)
- [Understanding IPv6 Support VDSL2 Interfaces on page 34](#)
- [Understanding MAC Limiting on Layer 3 Routing Interfaces on page 35](#)

Understanding Interface Logical Properties

Supported Platforms [SRX Series, vSRX](#)

The logical properties of an interface are the characteristics that do not apply to the physical interface or the wires connected to it. Logical properties include:

- Protocol families running on the interface (including any protocol-specific MTUs)
- IP address or addresses associated with the interface. A logical interface can be configured with an IPv6 address, IPv4 address, or both. The IP specification requires a unique address on every interface of each system attached to an IP network, so that traffic can be correctly routed. Individual hosts such as home computers must have a single IP address assigned. Devices must have a unique IP address for every interface.
- Virtual LAN (VLAN) tagging
- Any firewall filters or routing policies that are operating on the interface

Related Documentation

- [Understanding Interfaces on page 3](#)
- [Understanding Protocol Families on page 16](#)
- [Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 19](#)

- [Understanding Virtual LANs on page 47](#)

Understanding Protocol Families

Supported Platforms [SRX Series, vSRX](#)

A protocol family is a group of logical properties within an interface configuration. Protocol families include all the protocols that make up a protocol suite. To use a protocol within a particular suite, you must configure the entire protocol family as a logical property for an interface. The protocol families include common and not-so-common protocol suites.

This topic contains the following sections:

- [Common Protocol Suites on page 16](#)
- [Other Protocol Suites on page 16](#)

Common Protocol Suites

Junos OS protocol families include the following common protocol suites:

- **Inet**—Supports IP protocol traffic, including OSPF, BGP, and Internet Control Message Protocol (ICMP).
- **Inet6**—Supports IPv6 protocol traffic, including RIP for IPv6 (RIPng), IS-IS, and BGP.
- **ISO**—Supports IS-IS traffic.
- **MPLS**—Supports MPLS.



NOTE: Junos OS security features are flow-based—meaning the device sets up a flow to examine the traffic. Flow-based processing is not supported for ISO or MPLS protocol families.

Other Protocol Suites

In addition to the common protocol suites, Junos protocol families sometimes use the following protocol suites:

- **ccc**—Circuit cross-connect (CCC).
- **mlfr-uni-nni**—Multilink Frame Relay (MLFR) FRF.16 user-to-network network-to-network (UNI NNI).
- **mlfr-end-to-end**—Multilink Frame Relay end-to-end.
- **mlppp**—Multilink Point-to-Point Protocol.
- **tcc**—Translational cross-connect (TCC).
- **tnp**—Trivial Network Protocol. This Juniper Networks proprietary protocol provides communication between the Routing Engine and the device's packet forwarding

components. Junos OS automatically configures this protocol family on the device's internal interfaces only.

Related Documentation • [Understanding Interface Logical Properties on page 15](#)

Understanding IPv4 Addressing

Supported Platforms [SRX Series, vSRX](#)

IPv4 addresses are 32-bit numbers that are typically displayed in dotted decimal notation. A 32-bit address contains two primary parts: the network prefix and the host number.

All hosts within a single network share the same network address. Each host also has an address that uniquely identifies it. Depending on the scope of the network and the type of device, the address is either globally or locally unique. Devices that are visible to users outside the network (webservers, for example) must have a globally unique IP address. Devices that are visible only within the network must have locally unique IP addresses.

IP addresses are assigned by a central numbering authority called the Internet Assigned Numbers Authority (IANA). IANA ensures that addresses are globally unique where needed and has a large address space reserved for use by devices not visible outside their own networks.

This topic contains the following sections:

- [IPv4 Classful Addressing on page 17](#)
- [IPv4 Dotted Decimal Notation on page 18](#)
- [IPv4 Subnetting on page 18](#)
- [IPv4 Variable-Length Subnet Masks on page 19](#)

IPv4 Classful Addressing

To provide flexibility in the number of addresses distributed to networks of different sizes, 4-octet (32-bit) IP addresses were originally divided into three different categories or classes: class A, class B, and class C. Each address class specifies a different number of bits for its network prefix and host number:

- Class A addresses use only the first byte (octet) to specify the network prefix, leaving 3 bytes to define individual host numbers.
- Class B addresses use the first 2 bytes to specify the network prefix, leaving 2 bytes to define host addresses.
- Class C addresses use the first 3 bytes to specify the network prefix, leaving only the last byte to identify hosts.

In binary format, with an x representing each bit in the host number, the three address classes can be represented as follows:

```
00000000 xxxxxxxx xxxxxxxx xxxxxxxx (Class A)
00000000 00000000 xxxxxxxx xxxxxxxx (Class B)
00000000 00000000 00000000 xxxxxxxx (Class C)
```

Because each bit (x) in a host number can have a 0 or 1 value, each represents a power of 2. For example, if only 3 bits are available for specifying the host number, only the following host numbers are possible:

111 110 101 100 011 010 001 000

In each IP address class, the number of host-number bits raised to the power of 2 indicates how many host numbers can be created for a particular network prefix. Class A addresses have 2^{24} (or 16,777,216) possible host numbers, class B addresses have 2^{16} (or 65,536) host numbers, and class C addresses have 2^8 (or 256) possible host numbers.

IPv4 Dotted Decimal Notation

The 32-bit IPv4 addresses are most often expressed in dotted decimal notation, in which each octet (or byte) is treated as a separate number. Within an octet, the rightmost bit represents 2^0 (or 1), increasing to the left until the first bit in the octet is 2^7 (or 128).

Following are IP addresses in binary format and their dotted decimal equivalents:

11010000 01100010 11000000 10101010 = 208.98.192.170
 01110110 00001111 11110000 01010101 = 118.15.240.85
 00110011 11001100 00111100 00111011 = 51.204.60.59

IPv4 Subnetting

Because of the physical and architectural limitations on the size of networks, you often must break large networks into smaller subnetworks. Within a network, each wire or ring requires its own network number and identifying subnet address.

Figure 1 on page 18 shows two subnets in a network.

Figure 1: Subnets in a Network

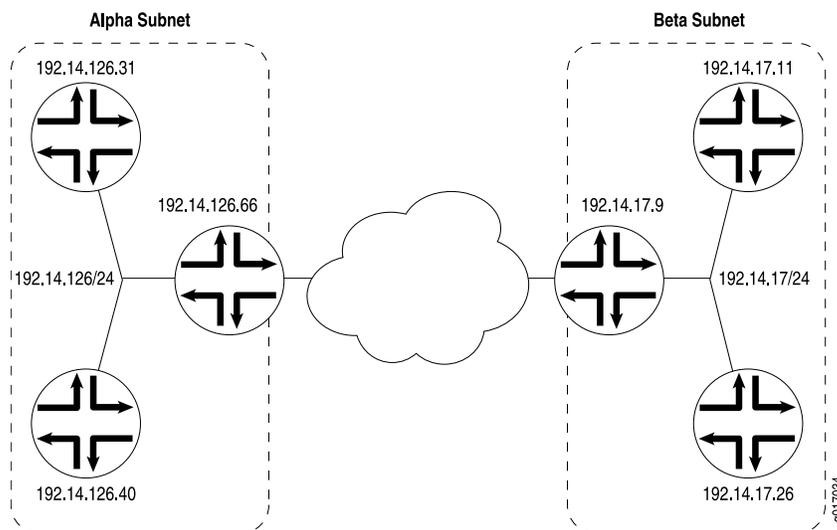


Figure 1 on page 18 shows three devices connected to one subnet and three more devices connected to a second subnet. Collectively, the six devices and two subnets make up the larger network. In this example, the network is assigned the network prefix **192.14.0.0**, a class B address. Each device has an IP address that falls within this network prefix.

In addition to sharing a network prefix (the first two octets), the devices on each subnet share a third octet. The third octet identifies the subnet. All devices on a subnet must have the same subnet address. In this case, the alpha subnet has the IP address **192.14.126.0** and the beta subnet has the IP address **192.14.17.0**.

The subnet address **192.14.17.0** can be represented as follows in binary notation:

```
11000000 . 00001110 . 00010001 . xxxxxxxx
```

Because the first 24 bits in the 32-bit address identify the subnet, the last 8 bits are not significant. To indicate the subnet, the address is written as **192.14.17.0/24** (or just **192.14.17/24**). The **/24** is the subnet mask (sometimes shown as **255.255.255.0**).

IPv4 Variable-Length Subnet Masks

Traditionally, subnets were divided by address class. Subnets had either 8, 16, or 24 significant bits, corresponding to 2^8 , 2^{16} , or 2^{24} possible hosts. As a result, an entire /16 subnet had to be allocated for a network that required only 400 addresses, wasting 65,136 ($2^{16} - 400 = 65,136$) addresses.

To help allocate address spaces more efficiently, variable-length subnet masks (VLSMs) were introduced. Using VLSM, network architects can allocate more precisely the number of addresses required for a particular subnet.

For example, suppose a network with the prefix **192.14.17/24** is divided into two smaller subnets, one consisting of 18 devices and the other of 46 devices.

To accommodate 18 devices, the first subnet must have 2^5 (32) host numbers. Having 5 bits assigned to the host number leaves 27 bits of the 32-bit address for the subnet. The IP address of the first subnet is therefore **192.14.17.128/27**, or the following in binary notation:

```
11000000 . 00001110 . 00010001 . 100xxxxx
```

The subnet mask includes 27 significant digits.

To create the second subnet of 46 devices, the network must accommodate 2^6 (64) host numbers. The IP address of the second subnet is **192.14.17.64/26**, or

```
11000000 . 00001110 . 00010001 . 01xxxxxx
```

By assigning address bits within the larger **/24** subnet mask, you create two smaller subnets that use the allocated address space more efficiently.

Related Documentation

- [Understanding Interface Logical Properties on page 15](#)
- [Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 19](#)

Understanding IPv6 Address Space, Addressing, Address Format, and Address Types

Supported Platforms [SRX Series, vSRX](#)

Understanding IP Version 6 (IPv6)

The ongoing expansive growth of the Internet and the need to provide IP addresses to accommodate it—to support increasing numbers of new users, computer networks, Internet-enabled devices, and new and improved applications for collaboration and communication—is escalating the emergent use of a new IP protocol. IPv6, with its robust architecture, was designed to satisfy these current and anticipated near future requirements.

IP version 4 (IPv4) is widely used throughout the world today for the Internet, intranets, and private networks. IPv6 builds upon the functionality and structure of IPv4 in the following ways:

- Provides a simplified and enhanced packet header to allow for more efficient routing.
- Improves support for mobile phones and other mobile computing devices.
- Enforces increased, mandatory data security through IPsec (which was originally designed for it).
- Provides more extensive quality-of-service (QoS) support.

IPv6 addresses consist of 128 bits, instead of 32 bits, and include a scope field that identifies the type of application suitable for the address. IPv6 does not support broadcast addresses, but instead uses multicast addresses for broadcast. In addition, IPv6 defines a new type of address called anycast.

Understanding IPv6 Address Types and How Junos OS for SRX Series Services Gateway Uses Them

IP version 6 (IPv6) includes the following types of addresses:

- Unicast

A unicast address specifies an identifier for a single interface to which packets are delivered. Under IPv6, the vast majority of Internet traffic is foreseen to be unicast, and it is for this reason that the largest assigned block of the IPv6 address space is dedicated to unicast addressing. Unicast addresses include all addresses other than loopback, multicast, link-local-unicast, and unspecified.

For SRX Series devices, the flow module supports the following kinds of IPv6 unicast packets:

- Pass-through unicast traffic, including traffic from and to virtual routers. The device transmits pass-through traffic according to its routing table.
- Host-inbound traffic from and to devices directly connected to SRX Series interfaces. For example, host-inbound traffic includes logging, routing protocol, and management types of traffic. The flow module sends these unicast packets to the Routing Engine and receives them from it. Traffic is processed by the Routing Engine instead of by the flow module, based on routing protocols defined for the Routing Engine.

The flow module supports all routing and management protocols that run on the Routing Engine. Some examples are OSPFv3, RIPng, TELNET, and SSH.

- Multicast

A multicast address specifies an identifier for a set of interfaces that typically belong to different nodes. It is identified by a value of 0xFF. IPv6 multicast addresses are distinguished from unicast addresses by the value of the high-order octet of the addresses.

The devices support only host-inbound and host-outbound multicast traffic. Host inbound traffic includes logging, routing protocols, management traffic, and so on.

- Anycast

An anycast address specifies an identifier for a set of interfaces that typically belong to different nodes. A packet with an anycast address is delivered to the nearest node, according to routing protocol rules.

There is no difference between anycast addresses and unicast addresses except for the subnet-router address. For an anycast subnet-router address, the low order bits, typically 64 or more, are zero. Anycast addresses are taken from the unicast address space.

The flow module treats anycast packets in the same way as it handles unicast packets. If an anycast packet is intended for the device, it is treated as host-inbound traffic, and it delivers it to the protocol stack which continues processing it.

IPv6 Address Scope

Unicast and multicast IPv6 addresses support address scoping, which identifies the application suitable for the address.

Unicast addresses support global address scope and two types of local address scope:

- Link-local unicast addresses—Used only on a single network link. The first 10 bits of the prefix identify the address as a link-local address. Link-local addresses cannot be used outside the link.
- Site-local unicast addresses—Used only within a site or intranet. A site consists of multiple network links. Site-local addresses identify nodes inside the intranet and cannot be used outside the site.

Multicast addresses support 16 different types of address scope, including node, link, site, organization, and global scope. A 4-bit field in the prefix identifies the address scope.

IPv6 Address Structure

Unicast addresses identify a single interface. Each unicast address consists of n bits for the prefix, and $128 - n$ bits for the interface ID.

Multicast addresses identify a set of interfaces. Each multicast address consists of the first 8 bits of all 1s, a 4-bit flags field, a 4-bit scope field, and a 112-bit group ID:

11111111 | f1gs | scop | group ID

The first octet of 1s identifies the address as a multicast address. The flags field identifies whether the multicast address is a well-known address or a transient multicast address.

The scope field identifies the scope of the multicast address. The 112-bit group ID identifies the multicast group.

Similar to multicast addresses, anycast addresses identify a set of interfaces. However, packets are sent to only one of the interfaces, not to all interfaces. Anycast addresses are allocated from the normal unicast address space and cannot be distinguished from a unicast address in format. Therefore, each member of an anycast group must be configured to recognize certain addresses as anycast addresses.

Understanding IPv6 Address Space, Addressing, and Address Types

Addressing is the area where most of the differences between IP version 4 (IPv4) and IPv6 exist, but the changes are largely about the ways in which addresses are implemented and used. IPv6 has a vastly larger address space than the impending exhausted IPv4 address space. IPv6 increases the size of the IP address from the 32 bits that compose an IPv4 address to 128 bits. Each extra bit given to an address doubles the size of the address space.

IPv4 has been extended using techniques such as Network Address Translation (NAT), which allows for ranges of private addresses to be represented by a single public address, and temporary address assignment. Although useful, these techniques fall short of the requirements of novel applications and environments such as emerging wireless technologies, always-on environments, and Internet-based consumer appliances.

In addition to the increased address space, IPv6 addresses differ from IPv4 addresses in the following ways:

- Includes a scope field that identifies the type of application that the address pertains to
- Does not support broadcast addresses, but instead uses multicast addresses to broadcast a packet
- Defines a new type of address, called anycast

Understanding IPv6 Address Format

All IPv6 addresses are 128 bits long, written as 8 sections of 16 bits each. They are expressed in hexadecimal representation, so the sections range from 0 to FFFF. Sections are delimited by colons, and leading zeroes in each section may be omitted. If two or more consecutive sections have all zeroes, they can be collapsed to a double colon.

IPv6 addresses consist of 8 groups of 16-bit hexadecimal values separated by colons (:). IPv6 addresses have the following format:

```
aaaa:aaaa:aaaa:aaaa:aaaa:aaaa:aaaa:aaaa
```

Each **aaaa** is a 16-bit hexadecimal value, and each **a** is a 4-bit hexadecimal value. Following is a sample IPv6 address:

```
3FFE:0000:0000:0001:0200:F8FF:FE75:50DF
```

You can omit the leading zeros of each 16-bit group, as follows:

```
3FFE:0:0:1:200:F8FF:FE75:50DF
```

You can compress 16-bit groups of zeros to double colons (::) as shown in the following example, but only once per address:

```
3FFE::1:200:F8FF:FE75:50DF
```

An IPv6 address prefix is a combination of an IPv6 prefix (address) and a prefix length. The prefix takes the form *ipv6-prefix/prefix-length* and represents a block of address space (or a network). The *ipv6-prefix* variable follows general IPv6 addressing rules. The */prefix-length* variable is a decimal value that indicates the number of contiguous, higher-order bits of the address that make up the network portion of the address. For example, 10FA:6604:8136:6502::/64 is a possible IPv6 prefix.

For more information on the text representation of IPv6 addresses and address prefixes, see RFC 4291, *IP Version 6 Addressing Architecture*.

Limitations

- On all branch SRX Series devices, changes in source AS and destination AS are not immediately reflected in exported flows.
- On all branch SRX Series devices, IPv6 traffic transiting over IPv4 based IP over IP tunnel (for example, IPv6-over-IPv4 using ip-x/x/x interface) is not supported.

Related Documentation

- *About the IPv6 Basic Packet Header*
- *Understanding IPv6 Packet Header Extensions*

Configuring the inet6 IPv6 Protocol Family

Supported Platforms SRX Series, vSRX

In configuration commands, the protocol family for IPv6 is named **inet6**. In the configuration hierarchy, instances of **inet6** are parallel to instances of **inet**, the protocol family for IPv4. In general, you configure **inet6** settings and specify IPv6 addresses in parallel to **inet** settings and IPv4 addresses.



NOTE: On SRX Series devices, on configuring identical IPs on a single interface, you will not see a warning message; instead, you will see a syslog message.

The following example shows the CLI commands you use to configure an IPv6 address for an interface:

```
[edit]
user@host# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 10.100.37.178/24;
    }
  }
}
```

```
[edit]
user@host# set interfaces ge-0/0/0 unit 0 family ?
Possible completions:
+ apply-groups          Groups from which to inherit configuration data
+ apply-groups-except  Don't inherit configuration data from these groups
> ccc                  Circuit cross-connect parameters
> ethernet-switching  Ethernet switching parameters
> inet                 IPv4 parameters
> inet6                IPv6 protocol parameters
> iso                  OSI ISO protocol parameters
> mpls                 MPLS protocol parameters
> tcc                  Translational cross-connect parameters
> vpls                 Virtual private LAN service parameters

[edit]
user@host# set interfaces ge-0/0/0 unit 0 family inet6 address 8d8d:8d01::1/64
user@host# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 10.100.37.178/24;
    }
    family inet6 {
      address 8d8d:8d01::1/64;
    }
  }
}
```

- Related Documentation**
- [Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 19](#)
 - [Enabling Flow-Based Processing for IPv6 Traffic on page 24](#)

Enabling Flow-Based Processing for IPv6 Traffic

Supported Platforms [SRX Series](#)

You have the following options for handling IPv6 traffic:

- Drop—Do not forward IPv6 packets. This is the default behavior.
- Packet-based forwarding—Do not create a session and process according to packet-based features only (includes firewall filters and class of service).
- Flow-based forwarding—Create a session and process according to packet-based features (including firewall filters and class of service) but also flow-based security features, such as screens and firewall security policy.

To enable flow-based processing for IPv6 traffic, modify the **mode** statement at the **[edit security forwarding-options family inet6]** hierarchy level:

```
security {
  forwarding-options {
    family {
      inet6 {
        mode flow-based;
      }
    }
  }
}
```

```

    }
  }
}

```

The following example shows the CLI commands you use to configure forwarding for IPv6 traffic:

```

[edit]
user@host# set security forwarding-options family inet6 mode ?
Possible completions:
  drop                Disable forwarding
  flow-based          Enable flow-based forwarding
  packet-based        Enable packet-based forwarding

[edit]
user@host# set security forwarding-options family inet6 mode flow-based
user@host# show security forwarding-options
family {
  inet6 {
    mode flow-based;
  }
}

```

If you change the forwarding option mode for IPv6, you might need to perform a reboot to initialize the configuration change. [Table 8 on page 25](#) summarizes device status upon configuration change.

Table 8: Device Status Upon Configuration Change

Configuration Change	Commit Warning	Reboot Required	Impact on Existing Traffic Before Reboot	Impact on New Traffic Before Reboot
Drop to flow-based	Yes	Yes	Dropped	Dropped
Drop to packet-based	No	No	Packet-based	Packet-based
Flow-based to packet-based	Yes	Yes	None	Flow sessions created
Flow-based to drop	Yes	Yes	None	Flow sessions created
Packet-based to flow-based	Yes	Yes	Packet-based	Packet-based
Packet-based to drop	No	No	Dropped	Dropped

- Related Documentation**
- [Understanding IPv6 Addressing](#)
 - [Configuring the inet6 IPv6 Protocol Family on page 23](#)

Configuring Flow Aggregation to Use Version 9 Flow Templates

Supported Platforms [SRX Series](#)

Use of version 9 allows you to define a flow record template suitable for IPv4 traffic, IPv6 traffic, or peer AS billing traffic. Templates and the fields included in the template are

transmitted to the collector periodically, and the collector need not be aware of the router configuration.



NOTE: Version 9 requires that you install a services PIC, such as the Adaptive Services PIC or the Multiservices PIC, in the device. On MX Series routers, the Multiservices DPC fulfills this requirement.

The following sections contain additional information:

- [Configuring the Traffic to Be Sampled on page 26](#)
- [Configuring the Version 9 Template Properties on page 26](#)
- [Restrictions on page 28](#)
- [Fields Included in Each Template Type on page 29](#)
- [inet Sampling Behavior on page 30](#)
- [Verification on page 31](#)
- [Examples: Configuring Version 9 Flow Templates on page 31](#)

Configuring the Traffic to Be Sampled

To specify sampling of IPv4, IPv6, or peer AS billing traffic, include the appropriate configuration of the **family** statement at the **[edit forwarding-options sampling input]** hierarchy level:

```
[edit forwarding-options sampling input]
family (inet ) {
  max-packets-per-second number;
  rate number;
  run-length number;
}
```

You can include **family inet**.



NOTE: If you specify sampling for peer AS billing traffic, the **family** statement supports only IPv4 and IPv6 traffic (**inet**). Peer AS billing traffic is enabled only at the global instance hierarchy level and is not available for per Packet Forwarding Engine instances.

Configuring the Version 9 Template Properties

To define the version 9 templates, include the following statements at the **[edit services flow-monitoring version9]** hierarchy level:

```
[edit services flow-monitoring version9]
template name {
  flow-active-timeout seconds;
  flow-inactive-timeout seconds;
  option-refresh-rate packets packets seconds seconds;
  template-refresh-rate packets packets seconds seconds;
```

```

(ipv4-template (Services) | ipv6-template (Services) | mpls-ipv4-template |
 mpls-template | peer-as-billing-template) {
  label-position [ positions ];
}
}

```

The following details apply to the configuration statements:

- You assign each template a unique name by including the **template name** statement.
- You then specify each template for the appropriate type of traffic by including the **ipv4-template**, **ipv6-template**, **inet-ipv4-template**, **inet-template**, or **peer-as-billing-template**.
- If the template is used for inet traffic, you can also specify up to three label positions for the inet header label data by including the **label-position** statement; the default values are [1 2 3].
- Within the template definition, you can optionally include values for the **flow-active-timeout** and **flow-inactive-timeout** statements. These statements have specific default and range values when they are used in template definitions; the default is 60 seconds and the range is from 10 through 600 seconds. Values you specify in template definitions override the global timeout values configured at the [edit forwarding-options sampling output flow-server] hierarchy level.



NOTE: In active flow monitoring, the flow-server records are exported after a time period that is a multiple of 60 seconds and greater than or equal to the configured active timeout value. For example, if the active timeout value is 90 seconds, the flow-server records are exported at 120-second intervals. If the active timeout value is 150 seconds, the flow-server records are exported at 180-second intervals, and so forth.

- You can also include settings for the **option-refresh-rate** and **template-refresh-rate** statements within a template definition. For both of these properties, you can include a timer value (in seconds) or a packet count (in number of packets). For the **seconds** option, the default value is 60 and the range is from 10 through 600. For the **packets** option, the default value is 4800 and the range is from 1 through 480,000.
- To filter IPV6 traffic on a media interface, the following configuration is supported:

```

interfaces interface-name {
  unit 0 {
    family inet {
      sampling {
        input;
        output;
      }
    }
  }
}

```

Restrictions

The following restrictions apply to version 9 templates:

- You cannot apply the two different types of flow aggregation configuration (flow-server version 5/8 and flow aggregation version 9) at the same time.
- Flow export based on an **inet-ipv4** template assumes that the IPv4 header follows the inet header. In the case of Layer 2 VPNs, the packet on the provider router (P router) would look like this:

inet | Layer 2 Header | IPv4

In this case, **inet-ipv4** flows are not created on the PIC, because the IPv4 header does not directly follow the inet header. Packets are dropped on the PIC and are accounted as parser errors.

- Outbound Routing Engine traffic is not sampled. A firewall filter is applied as output on the egress interface, which samples packets and exports the data. For transit traffic, egress sampling works correctly. For internal traffic, the next hop is installed in the Packet Forwarding Engine but sampled packets are not exported.
- Flows are created on the monitoring PIC only after the route record resynchronization operation is complete, which is 60 seconds after the PIC comes up. Any packets sent to the PIC would be dropped until the synchronization process is complete.

On all branch SRX Series devices, flow monitoring IPv6 version 9 has the following limitations:

- MPLS in not supported.
- User-defined version 9 templates are not supported.
- Routing Engine based flow monitoring version 9 is not supported.
- Flow monitoring and accounting are not supported in chassis cluster mode.
- Flow monitoring and accounting are not supported on an ae interface.
- J-Web for IPv6 sampled packets is not supported.
- SNMP queries for IPv6 sampled packets are not supported
- Flow monitoring can be configured in version 5, version 8, or version 9 export mode. Up to eight version 9 collectors are supported in export mode.
- Scope of accounting of IPv6 flow monitoring version 9 packets associated with pseudointerfaces (such as IRB, ML, LAG, VLAN, and GRE) is not supported.
- Creation of an SCTP session (parallel to TCP) between an exporter and a collector for gathering flow monitoring information is not supported.
- Maximum flow sessions that might be supported include:

- A device with 1-GB RAM, such as an SRX320 device, might support up to 15,000 flow monitoring sessions at a time.
- Routing Engine based flow monitoring V5 or V8 mode is mutually exclusive with inline flow monitoring V9.
- High-end SRX Series devices do not support multiple collectors like branch SRX Series devices. Only one V9 collector per IPv4 or IPv6 is supported
- Flow aggregation for V9 export is not supported.
- Only UDP over IPv4 or IPv6 protocol can be used as the transport protocol.
- Only the standard IPv4 or IPv6 template is supported for exporting flow monitoring records.
- User-defined or special templates are not supported for exporting flow monitoring records.
- Chassis cluster is supported without flow monitoring session synchronization.

Fields Included in Each Template Type

The following fields are common to all template types:

- Input interface
- Output interface
- Number of bytes
- Number of packets
- Flow start time
- Flow end time

The IPv4 template includes the following specific fields:

- IPv4 Source Address
- IPv4 Destination Address
- L4 Source Port
- L4 Destination Port
- IPv4 TOS
- IPv4 Protocol
- ICMP type and code
- TCP Flags
- IPv4 Next Hop Address

The IPv6 template includes the following specific fields:

- IPv6 Source Address and Mask
- IPv6 Destination Address and Mask
- L4 Source Port
- L4 Destination Port
- IPv6 TOS
- IPv6 Protocol
- TCP Flags
- IP Protocol Version
- IPv6 Next Hop Address
- Egress Interface Information
- Source Autonomous System (AS) number
- Destination AS number

The inet template includes the following specific fields:

- inet Label #1
- inet Label #2
- inet Label #3
- inet EXP Information
- FEC IP Address

The inet-IPv4 template includes all the fields found in the IPv4 and inet templates.

The peer AS billing template includes the following specific fields:

- IPV4 Class of Service (TOS)
- Ingress Interface
- BGP IPV4 Next Hop Address
- BGP Peer Destination AS Number

inet Sampling Behavior

This section describes the behavior when inet sampling is used on egress interfaces in various scenarios (label pop or swap) on provider routers (P routers).

1. You configure inet sampling on an egress interface on the P router and configure an inet flow aggregation template. The route action is label *pop* because penultimate hop popping (PHP) is enabled.

Previously, IPv4 packets (only) would have been sent to the PIC for sampling even though you configured inet sampling. No flows should be created, with the result that the parser fails.

With the current capability of applying inet templates, inet flows are created.

2. As in the first case, you configure inet sampling on an egress interface on the P router and configure an inet flow aggregation template. The route action is label swap and the swapped label is 0 (explicit null).

The resulting behavior is that inet packets are sent to the PIC. The flow being sampled corresponds to the label before the swap.

3. You configure a Layer 3 VPN network, in which a customer edge router (CE-1) sends traffic to a provider edge router (PE-A), through the P router, to a similar provider edge router (PE-B) and customer edge router (CE-2) on the remote end.

The resulting behavior is that you cannot sample inet packets on the PE-A to P router link.

Verification

To verify the configuration properties, you can use the **show services accounting aggregation template template-name name** operational mode command.

All other **show services accounting** commands also support version 9 templates, except for **show services accounting flow-detail** and **show services accounting aggregation aggregation-type**.

Examples: Configuring Version 9 Flow Templates

The following is a sample version 9 template configuration:

```
services {
  flow-monitoring {
    version9 {
      template ip-template {
        flow-active-timeout 20;
        flow-inactive-timeout 120;
        ipv4-template;
      }
      template inet-template-1 {
        inet-template {
          label-position [1 3 4];
        }
      }
      template inet-ipv4-template-1 {
        inet-ipv4-template {
          label-position [1 5 7];
        }
      }
      template peer-as-billing-template-1 {
        peer-as-billing-template;
      }
    }
  }
}
```

The following is a sample firewall filter configuration for inet traffic:

```
firewall {
  family inet {
    filter inet_sample {
      term default {
        then {
          accept;
          sample;
        }
      }
    }
  }
}
```

The following sample configuration applies the inet sampling filter on a networking interface and configures the AS PIC to accept both IPv4 and inet traffic:

```
inline-jflows {
  at-0/1/1 {
    unit 0 {
      family inet {
        filter {
          input inet_sample;
        }
      }
    }
  }
  sp-7/0/0 {
    unit 0 {
      family inet;
      family inet;
    }
  }
}
```

The following example applies the inet version 9 template to the sampling output and sends it to the AS PIC:

```
forwarding-options {
  sampling {
    input {
      family inet {
        rate 1;
      }
    }
    output {
      flow-active-timeout 60;
      flow-inactive-timeout 30;
      flow-server 1.2.3.4 {
        port 2055;
        version9 {
          template inet-ipv4-template-1;
        }
      }
    }
    inline-jflow sp-7/0/0 {
      source-address 1.1.1.1;
    }
  }
}
```

```

    }
  }
}

```

The following is a sample firewall filter configuration for the peer AS billing traffic:

```

firewall {
  family inet {
    filter peer-as-filter {
      term 0 {
        from {
          destination-class dcu-1;
          inline-jflow ge-2/1/0;
          forwarding-class class-1;
        }
        then count count_team_0;
      }
    }
    term 1 {
      from {
        destination-class dcu-2;
        inline-jflow ge-2/1/0;
        forwarding-class class-1;
      }
      then count count_team_1;
    }
    term 2 {
      from {
        destination-class dcu-3;
        inline-jflow ge-2/1/0;
        forwarding-class class-1;
      }
      then count count_team_2;
    }
  }
}

```

The following sample configuration applies the firewall filter as a filter attribute under the forwarding-options hierarchy for CoS-level data traffic usage information collection:

```

forwarding-options {
  family inet {
    filter output peer-as-filter;
  }
}

```

The following example applies the peer-as-billing version 9 template to enable sampling of traffic for billing purposes:

```

forwarding-options {
  sampling {
  }
  input {
    rate 1;
  }
  family inet {

```

```

output {
  flow-server 10.209.15.58 {
    port 300;
    version9 {
      template {
        peer-as;
      }
    }
    inline-jflow sp-5/2/0 {
      source-address 2.3.4.5;
    }
  }
}
family inet {
  filter {
    output peer-as-filter;
  }
}

```

Related Documentation

- [Understanding Interface Logical Properties on page 15](#)

Understanding IPv6 Support VDSL2 Interfaces

Supported Platforms SRX1500, SRX320, SRX340, SRX550M

The branch SRX Series devices support IPv6 on the following DSL encapsulations:

- ATM physical interface encapsulations
 - atm-pvc
 - ethernet-over-atm
- ATM logical interface encapsulations
 - atm-snap
 - atm-ppp-vc-mux
 - atm-nlpid
 - atm-cisco-nlpid
 - atm-ppp-llc
 - ether-over-atm-llc



NOTE: The encapsulation types atm-vc-mux and ppp-over-ether-over-atm-llc do not include IPv6 support.

To configure IPv6 addresses on DSL interfaces in ATM or PTM mode, include the family protocol type as **inet6** at the **[edit interfaces]** hierarchy level.

Related Documentation

- [Understanding Interface Logical Properties on page 15](#)

Understanding MAC Limiting on Layer 3 Routing Interfaces

Supported Platforms SRX1500, SRX300, SRX320, SRX340, vSRX

- [Overview on page 35](#)
- [Limitations on page 37](#)

Overview

The MAC limiting feature provides a mechanism for limiting MAC addresses on devices that are connected to a Layer 3 routed Gigabit Ethernet (GE), Fast Ethernet (FE), or 10 Gigabit Ethernet (XE) interface. With MAC filters, you can allow traffic with specific source MAC. Software-based MAC limiting is supported. MAC limiting is applicable only on interfaces with plain Ethernet or VLAN tagged encapsulation.

Both the physical interface level **source-address-filter** and logical interface level **accept-source-mac** configurations are supported on SRX300, SRX320, and SRX340 devices. The following considerations apply when you configure the **source-address-filter** and **accept-source-mac** statements:

- If only the logical level **accept-source-mac** statement is configured, traffic from only those configured MAC addresses will be allowed on the logical interface.
- If only the physical interface level **source-address-filter** statement is configured, the physical interface's *allowed* MAC addresses are also considered the *allowed* addresses for all the logical interfaces belonging to the physical interface. Incoming packets from any other source MAC addresses are dropped.
- If the physical interface level **source-address-filter** is configured under **gigether-options** (or **fastether-options**) and **accept-source-mac** is configured for one or more of its logical interfaces or VLANs, the allowed list of addresses is a combination of MAC addresses specified in both the statements. For logical interfaces and VLANs where the **accept-source-mac** statement is not configured, the physical interface's *allowed* list of addresses is considered.

You can configure an interface to receive packets from specific MAC addresses. To do this, specify the MAC addresses in the **source-address-filter** or **accept-source-mac** statements:

- **Logical level MAC filter configuration on an untagged interface**

```
ge-0/0/10 {
  unit 0 {
    accept-source-mac {
      mac-address 00:22:33:44:55:66;
      mac-address 00:26:88:e9:a3:01;
    }
  }
}
```

```

        family inet {
            address 60.60.60.1/24;
        }
    }
}

```

- Physical level MAC filter configuration on an untagged interface

```

ge-0/0/10 {
    gigeather-options {
        source-address-filter {
            00:55:55:55:55:66;
            00:26:88:e9:a3:01;
        }
    }
    unit 0 {
        family inet {
            address 60.60.60.1/24;
        }
    }
}

```

- Physical and logical level MAC filter configurations on a tagged interface

```

ge-0/0/10 {
    vlan-tagging;
    gigeather-options {
        source-address-filter {
            00:26:88:e9:a3:01;
        }
    }
    unit 0 {
        vlan-id 40;
        accept-source-mac {
            mac-address 00:22:33:44:55:66;
        }
        family inet {
            address 40.40.40.1/24;
        }
    }
    unit 1 {
        vlan-id 60;
        accept-source-mac {
            mac-address 00:55:55:55:55:66;
        }
        family inet {
            address 60.60.60.1/24;
        }
    }
}

```



NOTE: On untagged Gigabit Ethernet interfaces, you must not configure the `source-address-filter` and the `accept-source-mac` statements simultaneously. If these statements are configured for the same interfaces at the same time, an error message appears. However, in the case of tagged VLANs, both these statements can be configured simultaneously, if no identical MAC addresses are specified.

Limitations

The following limitations apply to MAC limiting support on Layer 3 routed GE, FE, or XE interfaces:

- You can configure only 32 MAC addresses per device.
- Only software-based MAC filtering is supported. Software-based MAC filtering impacts performance. The performance impact is proportional to the number of MAC addresses configured.
- MAC- based policer or rate limiting is not supported.
- You cannot configure broadcast or multicast address in the `source-address-filter` statement.
- MAC filtering is not supported on Aggregated Ethernet (AE), Fabric Ethernet, Point-to-Point Protocol over Ethernet (PPPoE), Routed VLAN interface (RVI), or VLAN interfaces.

MAC filtering is not supported on chassis clusters.

Related Documentation

- [Understanding Interface Logical Properties on page 15](#)

CHAPTER 3

Understanding Interface Physical Properties

- [Understanding Interface Physical Properties on page 39](#)
- [Understanding Bit Error Rate Testing on page 40](#)
- [Understanding Interface Clocking on page 41](#)
- [Understanding Frame Check Sequences on page 42](#)
- [MTU Default and Maximum Values on page 43](#)

Understanding Interface Physical Properties

Supported Platforms [SRX Series](#)

The physical properties of a network interface are the characteristics associated with the physical link that affect the transmission of either link-layer signals or the data across the links. Physical properties include clocking properties, transmission properties, such as the maximum transmission unit (MTU), and encapsulation methods, such as point-to-point and Frame Relay encapsulation.

The default property values for an interface are usually sufficient to successfully enable a bidirectional link. However, if you configure a set of physical properties on an interface, those same properties must be set on all adjacent interfaces to which a direct connection is made.

[Table 9 on page 39](#) summarizes some key physical properties of device interfaces.

Table 9: Interface Physical Properties

Physical Property	Description
<code>bert-error-rate</code>	Bit error rate (BER). The error rate specifies the number of bit errors in a particular bit error rate test (BERT) period required to generate a BERT error condition. See “Understanding Bit Error Rate Testing” on page 40 .
<code>bert-period</code>	Bit error rate test (BERT) time period over which bit errors are sampled. See “Understanding Bit Error Rate Testing” on page 40 .
<code>chap</code>	Challenge Handshake Authentication Protocol (CHAP). Specifying <code>chap</code> enables CHAP authentication on the interface. See “Understanding CHAP Authentication on a PPPoE Interface” on page 221 .

Table 9: Interface Physical Properties (*continued*)

Physical Property	Description
clocking	Clock source for the link. Clocking can be provided by the local system (internal) or a remote endpoint on the link (external). By default, all interfaces use the internal clocking mode. If an interface is configured to accept an external clock source, one adjacent interface must be configured to act as a clock source. Under this configuration, the interface operates in a loop timing mode, in which the clocking signal is unique for that individual network segment or loop. See “Understanding Interface Clocking” on page 41 .
description	A user-defined text description of the interface, often used to describe the interface's purpose.
disable	Administratively disables the interface.
encapsulation	Type of encapsulation on the interface. Common encapsulation types include PPP, Frame Relay, Cisco HDLC, and PPP over Ethernet (PPPoE). See “Understanding Physical Encapsulation on an Interface” on page 199 .
fcs	Frame check sequence (FCS). FCS is an error-detection scheme that appends parity bits to a digital signal and uses decoding algorithms that detect errors in the received digital signal.
mtu	Maximum transmission unit (MTU) size. MTU is the largest size packet or frame, specified in bytes or octets, that can be sent in a packet-based or frame-based network. The TCP uses MTU to determine the maximum size of each packet in any transmission. See “MTU Default and Maximum Values” on page 43 .
no-keepalives	Disabling of keepalive messages across a physical link. A keepalive message is sent between network devices to indicate that they are still active. Keepalives help determine whether the interface is operating correctly. Except for ATM-over-ADSL interfaces, all interfaces use keepalives by default.
pap	Password Authentication Protocol (PAP). Specifying pap enables PAP authentication on the interface. See “Understanding CHAP Authentication on a PPPoE Interface” on page 221 .
payload-scrambler	Scrambling of traffic transmitted out the interface. Payload scrambling randomizes the data payload of transmitted packets. Scrambling eliminates nonvariable bit patterns (strings of all 1s or all 0s) that generate link-layer errors across some physical links.

Related Documentation

- [Understanding Interfaces on page 3](#)
- [Understanding Bit Error Rate Testing on page 40](#)
- [Understanding Interface Clocking on page 41](#)
- [Understanding Frame Check Sequences on page 42](#)
- [MTU Default and Maximum Values on page 43](#)

Understanding Bit Error Rate Testing

Supported Platforms [SRX Series](#)

In telecommunication transmission, the bit error rate (BER) is the percentage of bits that have errors compared to the total number of bits received in a transmission, usually

expressed as 10 to a negative power. For example, a transmission with a BER of 10^{-6} received 1 errored bit in 1,000,000 bits transmitted. The BER indicates how often a packet or other data unit must be retransmitted because of an error. If the BER is too high, a slower data rate might improve the overall transmission time for a given amount of data if it reduces the BER and thereby lowers the number of resent packets.

A bit error rate test (BERT) is a procedure or device that measures the BER for a given transmission. You can configure a device to act as a BERT device by configuring the interface with a bit error rate and a testing period. When the interface receives a BERT request from a BER tester, it generates a response in a well-known BERT pattern. The initiating device checks the BERT-patterned response to determine the number of bit errors.

**Related
Documentation**

- [Understanding Interface Physical Properties on page 39](#)

Understanding Interface Clocking

Supported Platforms [SRX Series](#)

Clocking determines how individual routing nodes or entire networks sample transmitted data. As streams of information are received by a device in a network, a clock source specifies when to sample the data. In asynchronous networks, the clock source is derived locally, and synchronous networks use a central, external clock source. Interface clocking indicates whether the device uses asynchronous or synchronous clocking.



NOTE: Because truly synchronous networks are difficult to design and maintain, most synchronous networks are really plesiochronous networks. In a plesiochronous network, different timing regions are controlled by local clocks that are synchronized (with very narrow constraints). Such networks approach synchronicity and are generally known as synchronous networks.

Most networks are designed to operate as asynchronous networks. Each device generates its own clock signal, or devices use clocks from more than one clock source. The clocks within the network are not synchronized to a single clock source. By default, devices generate their own clock signals to send and receive traffic.

The system clock allows the device to sample (or detect) and transmit data being received and transmitted through its interfaces. Clocking enables the device to detect and transmit the 0s and 1s that make up digital traffic through the interface. Failure to detect the bits within a data flow results in dropped traffic.

Short-term fluctuations in the clock signal are known as *clock jitter*. Long-term variations in the signal are known as *clock wander*.

Asynchronous clocking can either derive the clock signal from the data stream or transmit the clocking signal explicitly.

This topic contains the following sections:

- [Data Stream Clocking on page 42](#)
- [Explicit Clocking Signal Transmission on page 42](#)

Data Stream Clocking

Common in T1 links, data stream clocking occurs when separate clock signals are not transmitted within the network. Instead, devices must extract the clock signal from the data stream. As bits are transmitted across the network, each bit has a time slot of 648 nanoseconds. Within a time slot, pulses are transmitted with alternating voltage peaks and drops. The receiving device uses the period of alternating voltages to determine the clock rate for the data stream.

Explicit Clocking Signal Transmission

Clock signals that are shared by hosts across a data link must be transmitted by one or both endpoints on the link. In a serial connection, for example, one host operates as a clock master and the other operates as a clock slave. The clock master internally generates a clock signal that is transmitted across the data link. The clock slave receives the clock signal and uses its period to determine when to sample data and how to transmit data across the link.

This type of clock signal controls only the connection on which it is active and is not visible to the rest of the network. An explicit clock signal does not control how other devices or even other interfaces on the same device sample or transmit data.

Related Documentation

- [Understanding Interface Physical Properties on page 39](#)

Understanding Frame Check Sequences

Supported Platforms [SRX Series](#)

All packets or frames within a network can be damaged by crosstalk or interference in the network's physical wires. The frame check sequence (FCS) is an extra field in each transmitted frame that can be analyzed to determine if errors have occurred. The FCS uses cyclic redundancy checks (CRCs), checksums, and two-dimensional parity bits to detect errors in the transmitted frames.

This topic contains the following sections:

- [Cyclic Redundancy Checks and Checksums on page 42](#)
- [Two-Dimensional Parity on page 43](#)

Cyclic Redundancy Checks and Checksums

On a link that uses CRCs for frame checking, the data source uses a predefined polynomial algorithm to calculate a CRC number from the data it is transmitting. The result is included in the FCS field of the frame and transmitted with the data. On the receiving end, the destination host performs the same calculation on the data it receives.

If the result of the second calculation matches the contents of the FCS field, the packet was sent and received without bit errors. If the values do not match, an FCS error is generated, the frame is discarded and the originating host is notified of the error.

Checksums function similarly to CRCs, but use a different algorithm.

Two-Dimensional Parity

On a link that uses two-dimensional parity bits for frame checking, the sending and receiving hosts examine each frame in the total packet transmission and create a parity byte that is evaluated to detect transmission errors.

For example, a host can create the parity byte for the following frame sequence by summing up each column (each bit position in the frame) and keeping only the least-significant bit:

Frame 1	0	1	0	1	0	0	1
Frame 2	1	1	0	1	0	0	1
Frame 3	1	0	1	1	1	1	0
Frame 4	0	0	0	1	1	1	0
Frame 5	0	1	1	0	1	0	0
Frame 6	1	0	1	1	1	1	1
Parity Byte	1	1	1	1	0	1	1

If the sum of the bit values in a bit position is even, the parity bit for the position is 0. If the sum is odd, the parity bit is 1. This method is called even parity. Matching parity bytes on the originating and receiving hosts indicate that the packet was received without error.

Related Documentation

- [Understanding Interface Physical Properties on page 39](#)

MTU Default and Maximum Values

Supported Platforms [SRX Series](#)

The MTU values are by default without any MTU configurations. If the MTU value is set, then the formula $IFF\ MTU\ (IP\ MTU) = IFD\ MTU\ (Media\ MTU) - L2\ Overhead$ is applicable. See [Table 10 on page 44](#) for default MTU values.



NOTE: For ATM MLPPP irrespective of UIFD MTU, the IP MTU is always 1500 because the IP MTU calculation is based on the LSQ interface. Even if you configure the LSQ family MTU, the IP MTU value cannot exceed 1504.

[Table 10 on page 44](#) lists MTU values for the SRX Series Services Gateways Physical Interface Modules (PIMs).

Table 10: MTU Values for the SRX Series Services Gateways PIMs

PIM	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP MTU (Bytes)
1-Port Gigabit Ethernet Small Form-Factor Pluggable (SFP) Mini-PIM	1514	9010	1500
1-Port Small Form-Factor Pluggable (SFP) Mini-PIM	1514	1518	1500
DOCSIS Mini-PIM	1504	1504	1500
Serial Mini-PIM	1504	2000	1500
T1/E1 Mini-PIM	1504	2000	1500
Dual CT1/E1 GPIM	1504	9000	1500
Quad CT1/E1 GPIM	1504	9000	1500
2-Port 10- Gigabit Ethernet XPIM	1514	9192	1500
16-Port Gigabit Ethernet XPIM	1514	9192	1500
24-Port Gigabit Ethernet XPIM	1514	9192	1500
ADSL2+ Mini-PIM (Encapsulation)			
atm-snap	1512	1512	1504
atm-vcmux	1512	1512	1512
atm-nlpid	1512	1512	1508
atm-cisco-nlpid	1512	1512	1510
ether-over-atm-llc	1512	1512	1488
atm-ppp-llc	1512	1512	1506
atm-ppp-vcmux	1512	1512	1510
atm-mlppp-llc	1512	1512	1500
ppp-over-ether-over-atm-llc	1512	1512	1480

Table 10: MTU Values for the SRX Series Services Gateways PIMs (*continued*)

PIM	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP MTU (Bytes)
VDSL- Mini-PIM AT mode (Encapsulation)			
atm-snap	1514	1514	1506
atm-vcmux	1514	1514	1514
atm-nlpid	1514	1514	1510
atm-cisco-nlpid	1514	1514	1512
ether-over-atm-llc	1514	1524	1490
atm-ppp-llc	1514	1514	1508
atm-ppp-vcmux	1514	1514	1512
atm-mlppp-llc	1514	1514	1500
ppp-over-ether-over-atm-llc	1514	1514	1482
VDSL- Mini-PIM PT mode			
	1514	1514	1500
G.SHDSL Mini-PIM AT mode (Encapsulation)			
atm-snap	4482	4482	4470
atm-vcmux	4482	4482	4470
atm-nlpid	4482	4482	4470
atm-cisco-nlpid	4482	4482	4470
ether-over-atm-llc	4482	4482	1500
atm-ppp-llc	4482	4482	4476
atm-ppp-vcmux	4482	4482	4480
atm-mlppp-llc	4482	4482	1500
ppp-over-ether-over-atm-llc	4482	4482	1492
G.SHDSL Mini-PIM PT mode			
	1514	1514	1500

- Related Documentation**
- [Understanding Interface Physical Properties on page 39](#)

CHAPTER 4

Configuring VLAN Tagging

- [Understanding Virtual LANs on page 47](#)
- [VLAN IDs and Ethernet Interface Types Supported on the SRX Series Devices on page 49](#)
- [Configuring VLAN Tagging on page 49](#)

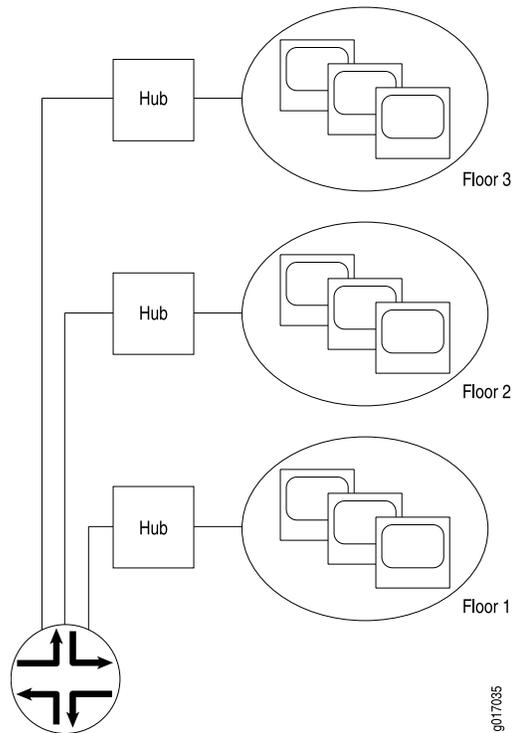
Understanding Virtual LANs

Supported Platforms [SRX Series, vSRX](#)

A LAN is a single broadcast domain. When traffic is broadcast, all hosts within the LAN receive the broadcast traffic. A LAN is determined by the physical connectivity of devices within the domain.

Within a traditional LAN, hosts are connected by a hub or repeater that propagates any incoming traffic throughout the network. Each host and its connecting hubs or repeaters make up a LAN segment. LAN segments are connected through switches and bridges to form the broadcast domain of the LAN. [Figure 2 on page 48](#) shows a typical LAN topology.

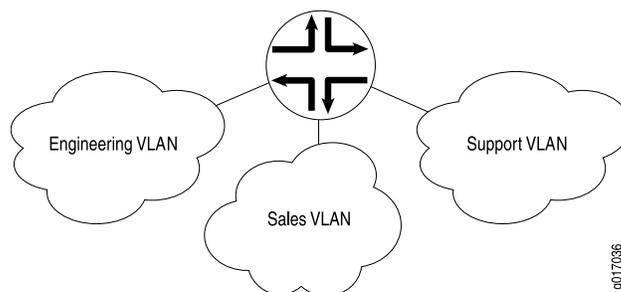
Figure 2: Typical LAN



Virtual LANs (VLANs) allow network architects to segment LANs into different broadcast domains based on logical groupings. Because the groupings are logical, the broadcast domains are not determined by the physical connectivity of the devices in the network. Hosts can be grouped according to a logical function, to limit the traffic broadcast within the VLAN to only the devices for which the traffic is intended.

Suppose a corporate network has three major organizations: engineering, sales, and support. Using VLAN tagging, hosts within each organization can be tagged with a different VLAN identifier. Traffic sent to the broadcast domain is then checked against the VLAN identifier and broadcast to only the devices in the appropriate VLAN. [Figure 3 on page 48](#) shows a typical VLAN topology.

Figure 3: Typical VLAN



- Related Documentation**
- [Understanding Interface Logical Properties on page 15](#)
 - [MPLS Feature Guide for Security Devices](#)

VLAN IDs and Ethernet Interface Types Supported on the SRX Series Devices

Supported Platforms SRX Series, vSRX

Table 11 on page 49 lists VLAN ID range by interface type supported on SRX Series devices:

Table 11: VLAN ID Range by Interface Type Supported on the SRX Series Devices

Interface Type	Interface Type VLAN ID Range
2-Port 10-Gigabit Ethernet	1 through 4094
10-Gigabit Ethernet	1 through 4094
16-Port Gigabit Ethernet	1 through 4094
24-Port Gigabit Ethernet	1 through 4094
Aggregated Ethernet for Fast Ethernet	1 through 1023
Aggregate Ethernet for Gigabit Ethernet	1 through 4094
Gigabit Ethernet	1 through 4094
Management and internal Ethernet interfaces	1 through 1023



NOTE: On SRX320, and SRX340 devices, on 1-GE SFP Mini-PIM, the VLAN ID 4093 falls under the reserved VLAN address range. Because of this, you will not be able to configure VLAN ID from this range.

Related Documentation

- [Understanding Interface Physical Properties on page 39](#)

Configuring VLAN Tagging

Supported Platforms SRX Series, vSRX

You can configure the branch SRX Series devices to receive and forward single-tag frames, dual-tag frames, or a mixture of single-tag and dual-tag frames.

See [Table 12 on page 49](#) for flexible VLANs.

Table 12: Flexible VLANs

Number of Tags	VLAN ID
0 (Untagged)	Native
1 (Tagged)	Single

Table 12: Flexible VLANs (*continued*)

Number of Tags	VLAN ID
2 (Dual tagged)	Dual

This topic includes the following sections:

- [Configuring Single-Tag Framing on page 50](#)
- [Configuring Dual Tagging on page 50](#)
- [Configuring Mixed Tagging on page 50](#)
- [Configuring Mixed Tagging Support for Untagged Packets on page 51](#)

Configuring Single-Tag Framing

To configure a device to receive and forward single-tag frames with 802.1Q VLAN tags, include the `vlan-tagging` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces interface-name]  
vlan-tagging;
```



NOTE: SRX Series high-end devices only support single-tag framing.

Configuring Dual Tagging

To configure the device to receive and forward dual-tag frames with 802.1Q VLAN tags, include the `flexible-vlan-tagging` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces interface-name]  
flexible-vlan-tagging;
```

Configuring Mixed Tagging

Mixed tagging is supported on ethernet interfaces of all branch SRX Series devices. Mixed tagging lets you configure two logical interfaces on the same Ethernet port, one with single-tag framing and one with dual-tag framing.

To configure mixed tagging, include the `flexible-vlan-tagging` statement at the `[edit interfaces ge-fpc/pic/port]` hierarchy level. You must also include the `vlan-tags` statement with `inner` and `outer` options or the `vlan-id` statement at the `[edit interfaces ge-fpc/pic/port unit logical-unit-number]` hierarchy level:

```
[edit interfaces ge-fpc/pic/port]  
flexible-vlan-tagging;  
unit logical-unit-number {  
  vlan-id number;  
  family family {  
    address address;  
  }  
}  
unit logical-unit-number {
```

```

vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
family family {
    address address;
}
}

```



NOTE: When you configure the physical interface MTU for mixed tagging, you must increase the MTU to 4 bytes more than the MTU value you would configure for a standard VLAN-tagged interface.

For example, if the MTU value is configured to be 1018 on a VLAN-tagged interface, then the MTU value on a flexible VLAN tagged interface must be 1022—4 bytes more. The additional 4 bytes accommodates the future addition of a stacked VLAN tag configuration on the same physical interface.

The following example configures mixed tagging. Dual-tag and single-tag logical interfaces are under the same physical interface:

```

[edit interfaces ge-0/2/0]
flexible-vlan-tagging;
unit 0 {
    vlan-id 232;
    family inet {
        address 10.66.1.2/30;
    }
}
unit 1 {
    vlan-tags outer 0x8100.222 inner 0x8100.221;
    family inet {
        address 10.66.1.2/30;
    }
}

```

Configuring Mixed Tagging Support for Untagged Packets

You can configure mixed tagging support for untagged packets on a port. Untagged packets are accepted on the same mixed VLAN-tagged port. To accept untagged packets, include the **native-vlan-id** statement and the **flexible-vlan-tagging** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```

[edit interfaces ge-fpc/pic/port]
flexible-vlan-tagging;
native-vlan-id number;

```

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

The following example configures untagged packets to be mapped to logical unit number 0:

```
[edit interfaces ge-0/2/0]
flexible-vlan-tagging;
native-vlan-id 232;
unit 0 {
  vlan-id 232;
  family inet {
    address 10.66.1.2/30;
  }
}
unit 1 {
  vlan-tags outer 0x8100.222 inner 0x8100.221;
  family inet {
    address 10.66.1.2/30;
  }
}
```

Related Documentation

- [Understanding Virtual LANs on page 47](#)

PART 2

Configuring DS1 Interfaces

- [Configuring DS1 Interfaces on page 55](#)
- [Configuring DS3 Interfaces on page 63](#)

CHAPTER 5

Configuring DS1 Interfaces

- [Understanding T1 and E1 Interfaces on page 55](#)
- [Example: Configuring a T1 Interface on page 58](#)
- [Example: Deleting a T1 Interface on page 60](#)

Understanding T1 and E1 Interfaces

Supported Platforms [SRX1500, SRX320, SRX340](#)

T1 and E1 are equivalent digital data transmission formats that carry DS1 signals. T1 and E1 lines can be interconnected for international use.

This topic contains the following sections:

- [T1 Overview on page 55](#)
- [E1 Overview on page 56](#)
- [T1 and E1 Signals on page 56](#)
- [Encoding on page 56](#)
- [T1 and E1 Framing on page 57](#)
- [T1 and E1 Loopback Signals on page 57](#)

T1 Overview

T1 is a digital data transmission medium capable of handling 24 simultaneous connections running at a combined 1.544 Mbps. T1 combines these 24 separate connections, called channels or time slots, onto a single link. T1 is also called DS1.

The T1 data stream is broken into frames. Each frame consists of a single framing bit and 24 8-bit channels, totaling 193 bits per T1 frame. Frames are transmitted 8,000 times per second, at a data transmission rate of 1.544 Mbps (8,000 x 193 = 1.544 Mbps).

As each frame is received and processed, the data in each 8-bit channel is maintained with the channel data from previous frames, enabling T1 traffic to be separated into 24 separate flows across a single medium. For example, in the following set of 4-channel frames (without a framing bit), the data in channel 1 consists of the first octet of each frame, the data in channel 2 consists of the second octet of each frame, and so on:

	Chan. 1	Chan. 2	Chan. 3	Chan. 4
Frame 1	[10001100]	[00110001]	[11111000]	[10101010]

```
Frame 2 [11100101] [01110110] [10001000] [11001010]
Frame 3 [00010100] [00101111] [11000001] [00000001]
```

E1 Overview

E1 is the European format for DS1 digital transmission. E1 links are similar to T1 links except that they carry signals at 2.048 Mbps. Each signal has 32 channels, and each channel transmits at 64 Kbps. E1 links have higher bandwidth than T1 links because they use all 8 bits of a channel. T1 links use 1 bit in each channel for overhead.

T1 and E1 Signals

T1 and E1 interfaces consist of two pairs of wires—a transmit data pair and a receive data pair. Clock signals, which determine when the transmitted data is sampled, are embedded in the T1 and E1 transmissions.

Typical digital signals operate by sending either zeros (0s) or ones (1s), which are usually represented by the absence or presence of a voltage on the line. The receiving device need only detect the presence of the voltage on the line at the particular sampling edge to determine whether the signal is 0 or 1. T1 and E1, however, use bipolar electrical pulses. Signals are represented by no voltage (0), positive voltage (1), or negative voltage (1). The bipolar signal allows T1 and E1 receivers to detect error conditions in the line, depending on the type of encoding that is being used.

Encoding

The following are common T1 and E1 encoding techniques:

- Alternate mark inversion (AMI)—T1 and E1
- Bipolar with 8-zero substitution (B8ZS)—T1 only
- High-density bipolar 3 code (HDB3)—E1 only

AMI Encoding

AMI encoding forces the 1s signals on a T1 or E1 line to alternate between positive and negative voltages for each successive 1 transmission, as in this sample data transmission:

```
1 1 0 1 0 1 0 1
+ - 0 + 0 - 0 +
```

When AMI encoding is used, a data transmission with a long sequence of 0s has no voltage transitions on the line. In this situation, devices have difficulty maintaining clock synchronization, because they rely on the voltage fluctuations to constantly synchronize with the transmitting clock. To counter this effect, the number of consecutive 0s in a data stream is restricted to 15. This restriction is called the 1s density requirement, because it requires a certain number of 1s for every 15 0s that are transmitted.

On an AMI-encoded line, two consecutive pulses of the same polarity—either positive or negative—are called a bipolar violation (BPV), which is generally flagged as an error.

B8ZS and HDB3 Encoding

Neither B8ZS nor HDB3 encoding restricts the number of 0s that can be transmitted on a line. Instead, these encoding methods detect sequences of 0s and substitute bit patterns for the sequences to provide the signal oscillations required to maintain timing on the link.

The B8ZS encoding method for T1 lines detects sequences of eight consecutive 0 transmissions and substitutes a pattern of two consecutive BPVs (11110000). Because the receiving end uses the same encoding, it detects the BPVs as 0s substitutions, and no BPV error is flagged. A single BPV, which does not match the 11110000 substitution bit sequence is likely to generate an error, depending on the configuration of the device.

The HDB3 encoding method for E1 lines detects sequences of four consecutive 0 transmissions and substitutes a single BPV (1100). Similar to B8ZS encoding, the receiving device detects the 0s substitutions and does not generate a BPV error.

T1 and E1 Framing

T1 interfaces use extended superframe (ESF). E1 interfaces use G.704 framing or G.704 with no CRC4 framing, or can be in unframed mode.

ESF Framing for T1

ESF extends the D4 superframe from 12 frames to 24 frames. By expanding the size of the superframe, ESF increases the number of bits in the superframe framing pattern from 12 to 24. The extra bits are used for frame synchronization, error detection, and maintenance communications through the facilities data link (FDL).

The ESF pattern for synchronization bits is 001011. Only the framing bits from frames 4, 8, 12, 16, 20, and 24 in the superframe sequence are used to create the synchronization pattern.

The framing bits from frames 2, 6, 10, 14, 18, and 22 are used to pass a CRC code for each superframe block. The CRC code verifies the integrity of the received superframe and detects bit errors with a CRC6 algorithm.

The framing bits for frames 1, 3, 5, 7, 9, 11, 13, 15, 17, 19, 21, and 23 are used for the data link channel. These 12 bits enable the operators at the network control center to query the remote equipment for information about the performance of the link.

T1 and E1 Loopback Signals

The control signal on a T1 or E1 link is the loopback signal. Using the loopback signal, the operators at the network control center can force the device at the remote end of a link to retransmit its received signals back onto the transmit path. The transmitting device can then verify that the received signals match the transmitted signals, to perform end-to-end checking on the link.

Two loopback signals are used to perform the end-to-end testing:

- The loop-up command signal sets the link into loopback mode, with the following command pattern:
...100001000010000100...
- The loop-down signal returns the link to its normal mode, with the following command pattern:
...100100100100100100...

While the link is in loopback mode, the operator can insert test equipment onto the line to test its operation.

Related Documentation

- [Example: Configuring a T1 Interface on page 58](#)

Example: Configuring a T1 Interface

Supported Platforms [SRX1500, SRX320, SRX340](#)

This example shows how to complete the initial configuration on a T1 interface.

- [Requirements on page 58](#)
- [Overview on page 58](#)
- [Configuration on page 58](#)
- [Verification on page 59](#)

Requirements

Before you begin, install a PIM, connect the interface cables to the ports, and power on the device. See the *Getting Started Guide* for your device.

Overview

This example describes the initial configuration that you must complete on each network interface. In this example, you configure the t1-1/0/0 interface as follows:

- You create the basic configuration for the new interface by setting the encapsulation type to ppp. You can enter additional values for physical interface properties as needed.
- You set the logical interface to 0. Note that the logical unit number can range from 0 through 16,384. You can enter additional values for properties you need to configure on the logical interface, such as logical encapsulation or protocol family.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter commit from configuration mode.

```
set interfaces t1-1/0/0 encapsulation ppp unit 0
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a T1 interface:

1. Create the interface.

```
[edit]
user@host# edit interfaces t1-1/0/0
```

2. Create the basic configuration for the new interface.

```
[edit interfaces t1-1/0/0]
user@host# set encapsulation ppp
```

3. Add logical interfaces.

```
[edit interfaces t1-1/0/0]
user@host# set unit 0
```

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

For brevity, this **show interfaces** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
...
t1-1/0/0 {
  encapsulation ppp;
  unit 0;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying the Link State of All Interfaces on page 59](#)
- [Verifying Interface Properties on page 60](#)

Verifying the Link State of All Interfaces

Purpose By using the ping tool on each peer address in the network, verify that all interfaces on the device are operational.

Action For each interface on the device:

1. In the J-Web interface, select **Troubleshoot>Ping Host**.

2. In the Remote Host box, type the address of the interface for which you want to verify the link state.
3. Click **Start**. The output appears on a separate page.

```

PING 10.10.10.10 : 56 data bytes
64 bytes from 10.10.10.10: icmp_seq=0 ttl=255 time=0.382 ms
64 bytes from 10.10.10.10: icmp_seq=1 ttl=255 time=0.266 ms

```

If the interface is operational, it generates an ICMP response. If this response is received, the round-trip time, in milliseconds, is listed in the time field.

Verifying Interface Properties

Purpose Verify that the interface properties are correct.

Action From the operational mode, enter the **show interfaces detail** command.

The output shows a summary of interface information. Verify the following information:

- The physical interface is Enabled. If the interface is shown as Disabled, do one of the following:
 - In the CLI configuration editor, delete the **disable** statement at the [edit interfaces t1-1/0/0] level of the configuration hierarchy.
 - In the J-Web configuration editor, clear the **Disable** check box on the Interfaces > t1-1/0/0 page.
- The physical link is Up. A link state of Down indicates a problem with the interface module, interface port, or physical connection (link-layer errors).
- The Last Flapped time is an expected value. It indicates the last time the physical interface became unavailable and then available again. Unexpected flapping indicates likely link-layer errors.
- The traffic statistics reflect expected input and output rates. Verify that the number of input and output bytes and packets matches expected throughput for the physical interface. To clear the statistics and see only new changes, use the **clear interfaces statistics t1-1/0/0** command.

- Related Documentation**
- [Understanding T1 and E1 Interfaces on page 55](#)
 - [Example: Deleting a T1 Interface on page 60](#)

Example: Deleting a T1 Interface

Supported Platforms SRX1500, SRX320, SRX340

This example shows how to delete a T1 interface.

- [Requirements on page 61](#)
- [Overview on page 61](#)
- [Configuration on page 61](#)
- [Verification on page 61](#)

Requirements

No special configuration beyond device initialization is required before configuring an interface.

Overview

In this example, you delete the t1-1/0/0 interface.



NOTE: Performing this action removes the interface from the software configuration and disables it. Network interfaces remain physically present, and their identifiers continue to appear on the J-Web pages.

Configuration

Step-by-Step Procedure

To delete a T1 interface:

1. Specify the interface you want to delete.

```
[edit interfaces]  
user@host# delete t1-1/0/0
```
2. If you are done configuring the device, commit the configuration.

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

Verification

To verify the configuration is working properly, enter the **show interfaces** command.

Related Documentation

- [Understanding T1 and E1 Interfaces on page 55](#)
- [Example: Configuring a T1 Interface on page 58](#)

CHAPTER 6

Configuring DS3 Interfaces

- [Understanding T3 and E3 Interfaces on page 63](#)
- [Example: Configuring a T3 Interface on page 68](#)
- [Example: Deleting a T3 Interface on page 70](#)

Understanding T3 and E3 Interfaces

Supported Platforms SRX1500

T3 is a high-speed data-transmission medium formed by multiplexing 28 DS1 signals into seven separate DS2 signals, and combining the DS2 signals into a single DS3 signal. T3 links operate at 43.736 Mbps. T3 is also called DS3.

E3 is the equivalent European transmission format. E3 links are similar to T3 (DS3) links, but carry signals at 34.368 Mbps. Each signal has 16 E1 channels, and each channel transmits at 2.048 Mbps. E3 links use all 8 bits of a channel, whereas T3 links use 1 bit in each channel for overhead.

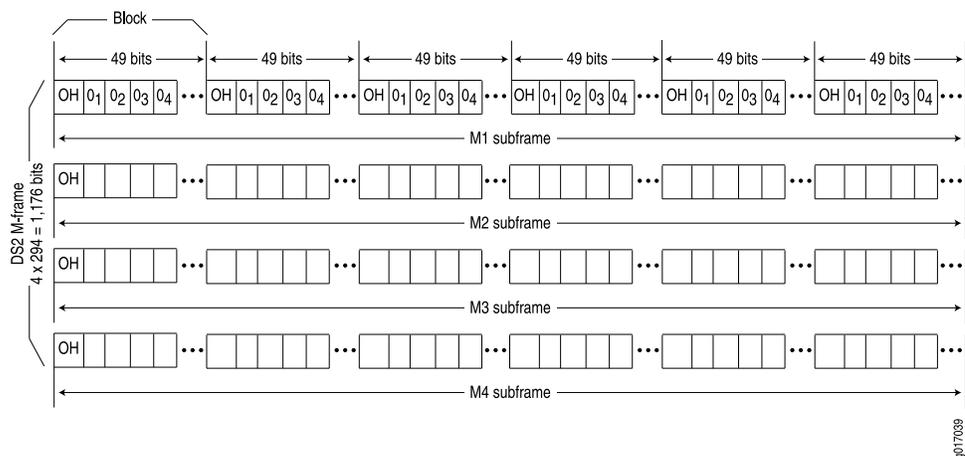
- [Multiplexing DS1 Signals on page 63](#)
- [DS2 Bit Stuffing on page 64](#)
- [DS3 Framing on page 64](#)

Multiplexing DS1 Signals

Four DS1 signals combine to form a single DS2 signal. The four DS1 signals form a single DS2 M-frame, which includes subframes M1 through M4. Each subframe has six 49-bit blocks, for a total of 294 bits per subframe. The first bit in each block is a DS2 overhead (OH) bit. The remaining 48 bits are DS1 information bits.

[Figure 4 on page 64](#) shows the DS2 M-frame format.

Figure 4: DS2 M-Frame Format



The four DS2 subframes are not four DS1 channels. Instead, the DS1 data bits within the subframes are formed by data interleaved from the DS1 channels. The O_n values designate time slots devoted to DS1 inputs as part of the bit-by-bit interleaving process. After every 48 DS1 information bits (12 bits from each signal), a DS2 OH bit is inserted to indicate the start of a subframe.

DS2 Bit Stuffing

Because the four DS1 signals are asynchronous signals, they might operate at different line rates. To synchronize the asynchronous streams, the multiplexers on the line use bit stuffing.

A DS2 connection requires a nominal transmit rate of 6.304 Mbps. However, because multiplexers increase the overall output rate to the intermediate rate of 6.312 Mbps, the output rate is higher than individual input rates on DS1 signals. The extra bandwidth is used to stuff the incoming DS1 signals with extra bits until the output rate of each signal equals the increased intermediate rate. These stuffed bits are inserted at fixed locations in the DS2 M-frame. When DS2 frames are received and the signal is demultiplexed, the stuffing bits are identified and removed.

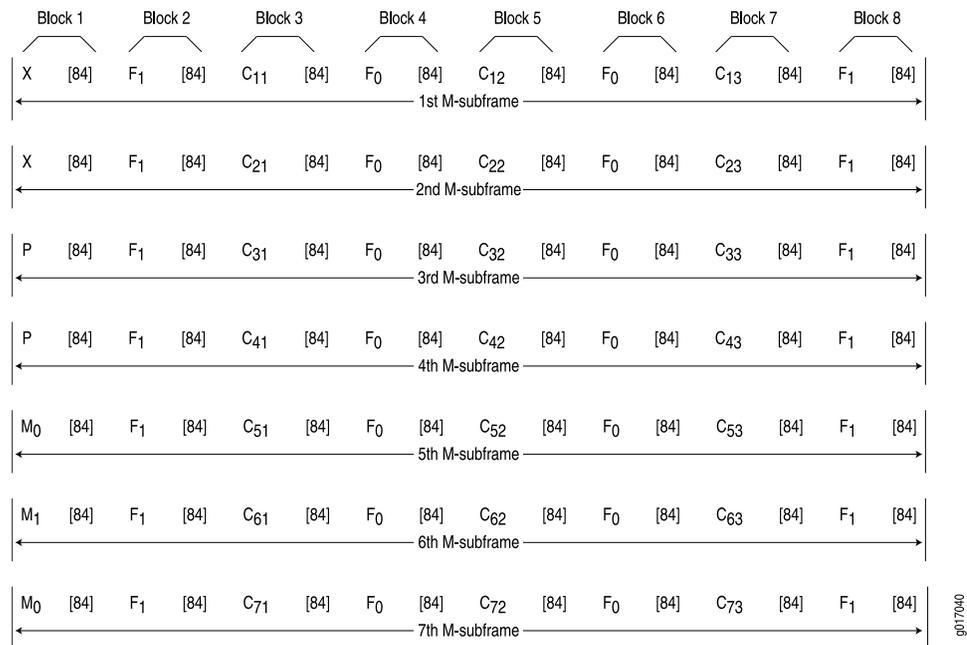
DS3 Framing

A set of four DS1 signals is multiplexed into seven DS2 signals, which are multiplexed into a single DS3 signal. The multiplexing occurs just as with DS1-to-DS2 multiplexing. The resulting DS3 signal uses either the standard M13 asynchronous framing format or the C-bit parity framing format. Although the two framing formats differ in their use of control and message bits, the basic frame structures are identical. The DS3 frame structures are shown in [Figure 5 on page 65](#) and [Figure 6 on page 66](#).

M13 Asynchronous Framing

A DS3 M-frame includes seven subframes, formed by DS2 data bits interleaved from the seven multiplexed DS2 signals. Each subframe has eight 85-bit blocks—a DS3 OH bit plus 84 data bits. The meaning of an OH bit depends on the block it precedes. Standard DS3 M13 asynchronous framing format is shown in [Figure 5 on page 65](#).

Figure 5: DS3 M13 Frame Format



A DS3 M13 M-frame contains the following types of OH bits:

- Framing bits (F-bits)—Make up a frame alignment signal that synchronizes DS3 subframes. Each DS3 frame contains 28 F-bits (4 bits per subframe). F-bits are located at the beginning of blocks 2, 4, 6, and 8 of each subframe. When combined, the frame alignment pattern for each subframe is 1001. The pattern can be examined to detect bit errors in the transmission.
- Multiframe bits (M-bits)—Make up a multiframe alignment signal that synchronizes the M-frames in a DS3 signal. Each DS3 frame contains 3 M-bits, which are located at the beginning of subframes 5, 6, and 7. When combined, the multiframe alignment pattern for each M-frame is 010.
- Bit stuffing control bits (C-bits)—Serve as bit stuffing indicators for each DS2 input. For example, C₁₁, C₁₂, and C₁₃ are indicators for DS2 input 1. Their values indicate whether DS3 bit stuffing has occurred at the multiplexer. If the three C-bits in a subframe are all 0s, no stuffing was performed for the DS2 input. If the three C-bits are all 1s, stuffing was performed.
- Message bits (X-bits)—Used by DS3 transmitters to embed asynchronous in-service messages in the data transmission. Each DS3 frame contains 2 X-bits, which are located at the beginning of subframes 1 and 2. Within an DS3 M-frame, both X-bits must be identical.
- Parity bits (P-bits)—Compute parity over all but 1 bit of the M-frame. (The first X-bit is not included.) Each DS3 frame contains 2 P-bits, which are located at the beginning of subframes 3 and 4. Both P-bits must be identical.

If the previous DS3 frame contained an odd number of 1s, both P-bits are set to 1. If the previous DS3 contained an even number of 1s, both P-bits are set to 0. If, on the receiving

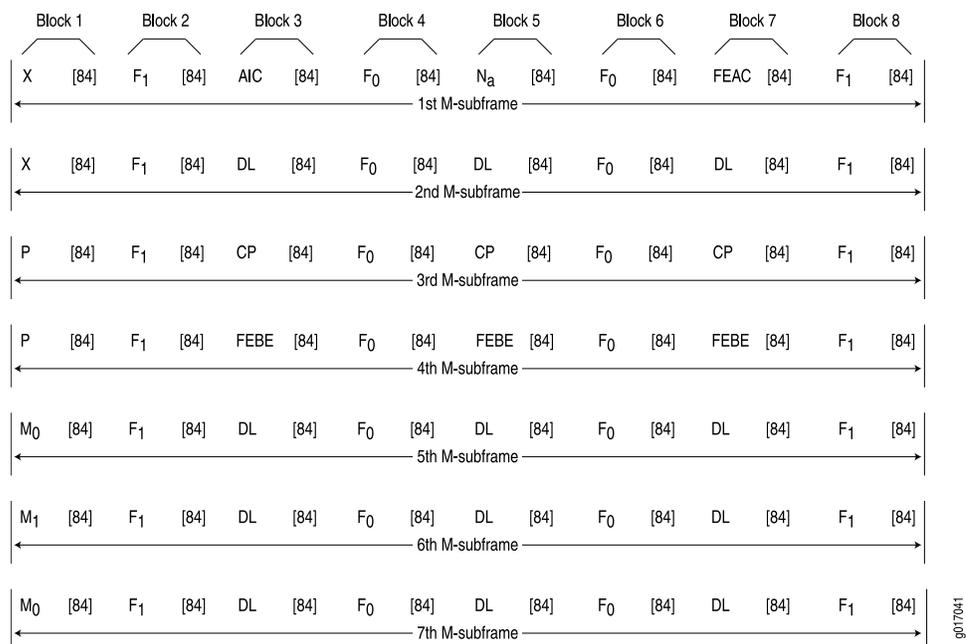
side, the number of 1s for a given frame does not match the P-bits in the following frame, it indicates one or more bit errors in the transmission.

C-Bit Parity Framing

In M13 framing, every C-bit in a DS3 frame is used for bit stuffing. However, because multiplexers first use bit stuffing when multiplexing DS1 signals into DS2 signals, the incoming DS2 signals are already synchronized. Therefore, the bit stuffing that occurs when DS2 signals are multiplexed is redundant.

C-bit parity framing format redefines the function of C-bits and X-bits, using them to monitor end-to-end path performance and provide in-band data links. The C-bit parity framing structure is shown in Figure 6 on page 66.

Figure 6: DS3 C-Bit Parity Framing



In C-bit parity framing, the X-bits transmit error conditions from the far end of the link to the near end. If no error conditions exist, both X-bits are set to 1. If an out-of-frame (OOF) or alarm indication signal (AIS) error is detected, both X-bits are set to 0 in the upstream direction for 1 second to notify the other end of the link about the condition.

The C-bits that control bit stuffing in M13 frames are typically used in the following ways by C-bit parity framing:

- Application identification channel (AIC)—The first C-bit in the first subframe identifies the type of DS3 framing used. A value of 1 indicates that C-bit parity framing is in use.
- N_a—A reserved network application bit.
- Far-end alarm and control (FEAC) channel—The third C-bit in the first subframe is used for the FEAC channel. In normal transmissions, the FEAC C-bit transmits all 1s.

When an alarm condition is present, the FEAC C-bit transmits a code word in the format **0xxxxxx 11111111**, in which *x* can be either 1 or 0. Bits are transmitted from right to left.

[Table 13 on page 67](#) lists some C-bit code words and the alarm or status condition indicated.

Table 13: FEAC C-Bit Condition Indicators

Alarm or Status Condition	C-Bit Code Word
DS3 equipment failure requires immediate attention.	00110010 11111111
DS3 equipment failure occurred—such as suspended, not activated, or unavailable service—that is non-service-affecting.	00011110 11111111
DS3 loss of signal.	00011100 11111111
DS3 out of frame.	00000000 11111111
DS3 alarm indication signal (AIS) received.	00101100 11111111
DS3 idle received.	00110100 11111111
Common equipment failure occurred that is non-service-affecting.	00011101 11111111
Multiple DS1 loss of signal.	00101010 11111111
DS1 equipment failure occurred that requires immediate attention.	00001010 11111111
DS1 equipment failure occurred that is non-service-affecting.	00000110 11111111
Single DS1 loss of signal.	00111100 11111111

- **Data links**—The 12 C-bits in subframes 2, 5, 6, and 7 are data link (DL) bits for applications and terminal-to-terminal path maintenance.
- **DS3 parity**—The 3 C-bits in the third subframe are DS3 parity C-bits (also called CP-bits). When a DS3 frame is transmitted, the sending device sets the CP-bits to the same value as the P-bits. When the receiving device processes the frame, it calculates the parity of the M-frame and compares this value to the parity in the CP-bits of the following M-frame. If no bit errors have occurred, the two values are typically the same.
- **Far-end block errors (FEBEs)**—The 3 C-bits in the fourth subframe make up the far-end block error (FEBE) bits. If a framing or parity error is detected in an incoming M-frame (via the CP-bits), the receiving device generates a C-bit parity error and sends an error notification to the transmitting (far-end) device. If an error is generated, the FEBE bits are set to 000. If no error occurred, the bits are set to 111.

Related Documentation

- [Example: Configuring a T3 Interface on page 68](#)
- [Example: Deleting a T3 Interface on page 70](#)

Example: Configuring a T3 Interface

Supported Platforms **SRX1500**

This example shows how to complete the initial configuration on a T3 interface.

- [Requirements on page 68](#)
- [Overview on page 68](#)
- [Configuration on page 68](#)
- [Verification on page 69](#)

Requirements

Before you begin, install a PIM, connect the interface cables to the ports, and power on the device. See the *Getting Started Guide* for your device.

Overview

This example describes the initial configuration that you must complete on each network interface. In this example, you configure the t3-1/0/0 interface as follows:

- You create the basic configuration for the new interface by setting the encapsulation type to ppp. You can enter additional values for physical interface properties as needed.
- You set the logical interface to 0. Note that the logical unit number can range from 0 to 16,384. You can enter additional values for properties you need to configure on the logical interface, such as logical encapsulation or protocol family.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter commit from configuration mode.

```
set interfaces t3-1/0/0 encapsulation ppp unit 0
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a T3 interface:

1. Create the interface.

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

2. Create the basic configuration for the new interface.

```
[edit interfaces t3-1/0/0]  
user@host# set encapsulation ppp
```

3. Add logical interfaces.

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

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

For brevity, this **show interfaces** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
...
t3-1/0/0 {
  encapsulation ppp;
  unit 0;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying the Link State of All Interfaces on page 69](#)
- [Verifying Interface Properties on page 69](#)

Verifying the Link State of All Interfaces

Purpose By using the ping tool on each peer address in the network, verify that all interfaces on the device are operational.

Action For each interface on the device:

1. In the J-Web interface, select **Troubleshoot>Ping Host**.
2. In the Remote Host box, type the address of the interface for which you want to verify the link state.
3. Click **Start**. The output appears on a separate page.

```
PING 10.10.10.10 : 56 data bytes
64 bytes from 10.10.10.10: icmp_seq=0 ttl=255 time=0.382 ms
64 bytes from 10.10.10.10: icmp_seq=1 ttl=255 time=0.266 ms
```

If the interface is operational, it generates an ICMP response. If this response is received, the round-trip time in milliseconds is listed in the time field.

Verifying Interface Properties

Purpose Verify that the interface properties are correct.

Action From the operational mode, enter the **show interfaces detail** command.

The output shows a summary of interface information. Verify the following information:

- The physical interface is Enabled. If the interface is shown as Disabled, do one of the following:
 - In the CLI configuration editor, delete the **disable** statement at the [edit interfaces t3-1/0/0] level of the configuration hierarchy.
 - In the J-Web configuration editor, clear the **Disable** check box on the Interfaces> t3-1/0/0 page.
- The physical link is Up. A link state of Down indicates a problem with the interface module, interface port, or physical connection (link-layer errors).
- The Last Flapped time is an expected value. It indicates the last time the physical interface became unavailable and then available again. Unexpected flapping indicates likely link-layer errors.
- The traffic statistics reflect expected input and output rates. Verify that the number of input and output bytes and packets matches expected throughput for the physical interface. To clear the statistics and see only new changes, use the **clear interfaces statistics t3-1/0/0** command.

- Related Documentation**
- [Understanding T3 and E3 Interfaces on page 63](#)
 - [Example: Deleting a T3 Interface on page 70](#)

Example: Deleting a T3 Interface

Supported Platforms SRX1500

This example shows how to delete a T3 interface.

- [Requirements on page 70](#)
- [Overview on page 70](#)
- [Configuration on page 71](#)
- [Verification on page 71](#)

Requirements

No special configuration beyond device initialization is required before configuring an interface.

Overview

In this example, you delete the t3-1/0/0 interface.



NOTE: Performing this action removes the interface from the software configuration and disables it. Network interfaces remain physically present, and their identifiers continue to appear on the J-Web pages.

Configuration

Step-by-Step Procedure

To delete a T3 interface:

1. Specify the interface you want to delete.

```
[edit interfaces]  
user@host# delete t3-1/0/0
```

2. If you are done configuring the device, commit the configuration.

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

Verification

To verify the configuration is working properly, enter the **show interfaces** command.

Related Documentation

- [Understanding T3 and E3 Interfaces on page 63](#)
- [Example: Configuring a T3 Interface on page 68](#)

PART 3

Configuring DSL Interfaces

- [Configuring VDSL2 Interfaces on page 75](#)

CHAPTER 7

Configuring VDSL2 Interfaces

- [VDSL2 Interface Technology Overview on page 75](#)
- [VDSL2 Network Deployment Topology on page 76](#)
- [VDSL2 Interface Support on SRX Series Devices on page 77](#)
- [Example: Configuring VDSL2 Interfaces \(Basic\) on page 81](#)
- [Example: Configuring VDSL2 Interfaces \(Detail\) on page 87](#)
- [Upgrading the VDSL PIC Firmware on page 112](#)

VDSL2 Interface Technology Overview

Supported Platforms [SRX320, SRX340](#)

Very-high-bit-rate digital subscriber line (VDSL) technology is part of the xDSL family of modem technologies that provide faster data transmission over a single flat untwisted or twisted pair of copper wires. The VDSL lines connect service provider networks and customer sites to provide high bandwidth applications (triple-play services) such as high-speed Internet access, telephone services like VoIP, high-definition TV (HDTV), and interactive gaming services over a single connection.

VDSL2 is an enhancement to G.993.1 (VDSL) and permits the transmission of asymmetric (half-duplex) and symmetric (full-duplex) aggregate data rates up to 100 Mbps on short copper loops using a bandwidth up to 30 MHz. The VDSL2 technology is based on the ITU-T G.993.2 (VDSL2) standard, which is the International Telecommunication Union standard describing a data transmission method for VDSL2 transceivers.

The VDSL2 uses discrete multitone (DMT) modulation. DMT is a method of separating a digital subscriber line signal so that the usable frequency range is separated into 256 frequency bands (or channels) of 4.3125 KHz each. The DMT uses the Fast Fourier Transform (FFT) algorithm for demodulation or modulation for increased speed.

VDSL2 interface supports Packet Transfer Mode (PTM). The PTM mode transports packets (IP, PPP, Ethernet, MPLS, and so on) over DSL links as an alternative to using Asynchronous Transfer Mode (ATM). PTM is based on the Ethernet in the First Mile (EFM) IEEE802.3ah standard.

VDSL2 provides backward compatibility with ADSL, ADSL2, and ADSL2+ because this technology is based on both the VDSL1-DMT and ADSL2/ADSL2+ recommendations.

VDSL2 Vectoring Overview

Vectoring is a transmission method that employs the coordination of line signals that reduce crosstalk levels and improve performance. It is based on the concept of noise cancellation, like noise-cancelling headphones. The ITU-T G.993.5 standard, "Self-FEXT Cancellation (Vectoring) for Use with VDSL2 Transceivers," also known as G.vector, describes vectoring for VDSL2.

The scope of Recommendation ITU-T G.993.5 is specifically limited to the self-FEXT (far-end crosstalk) cancellation in the downstream and upstream directions. The FEXT generated by a group of near-end transceivers and interfering with the far-end transceivers of that same group is canceled. This cancellation takes place between VDSL2 transceivers, not necessarily of the same profile.

Related Documentation

- [VDSL2 Network Deployment Topology on page 76](#)
- [VDSL2 Interface Support on SRX Series Devices on page 77](#)
- [Example: Configuring VDSL2 Interfaces \(Basic\) on page 81](#)
- [Example: Configuring VDSL2 Interfaces \(Detail\) on page 87](#)

VDSL2 Network Deployment Topology

Supported Platforms [SRX320, SRX340](#)

In standard telephone cables of copper wires, voice signals use only a fraction of the available bandwidth. Like any other DSL technology, the VDSL2 technology utilizes the remaining capacity to carry the data and multimedia on the wire without interrupting the line's ability to carry voice signals.

This example depicts the typical VDSL2 network topology deployed using SRX Series Services Gateways.

A VDSL2 link between network devices is set up as follows:

1. Connect an end-user device such as a LAN, hub, or PC through an Ethernet interface to the customer premises equipment (CPE) (for example, an SRX Series device).
2. Connect the CPE to a DSLAM.
3. The VDSL2 interface uses either Gigabit Ethernet or fiber as second mile to connect to the Broadband Remote Access Server (B-RAS) as shown in [Figure 7 on page 77](#).
4. The ADSL interface uses either Gigabit Ethernet (in case of IP DSLAM] as the "second mile" to connect to the B-RAS or OC3/DS3 ATM as the second mile to connect the B-RAS as shown in [Figure 8 on page 77](#).



NOTE: The VDSL2 technology is backward compatible with ADSL. VDSL2 provides an ADSL interface in an ATM DSLAM topology and provides a VDSL2 interface in an IP or VDSL DSLAM topology.

The DSLAM accepts connections from many customers and aggregates them to a single, high-capacity connection to the Internet.

Figure 7 on page 77 shows a typical VDSL2 network topology.

Figure 7: Typical VDSL2 End-to-End Connectivity and Topology Diagram

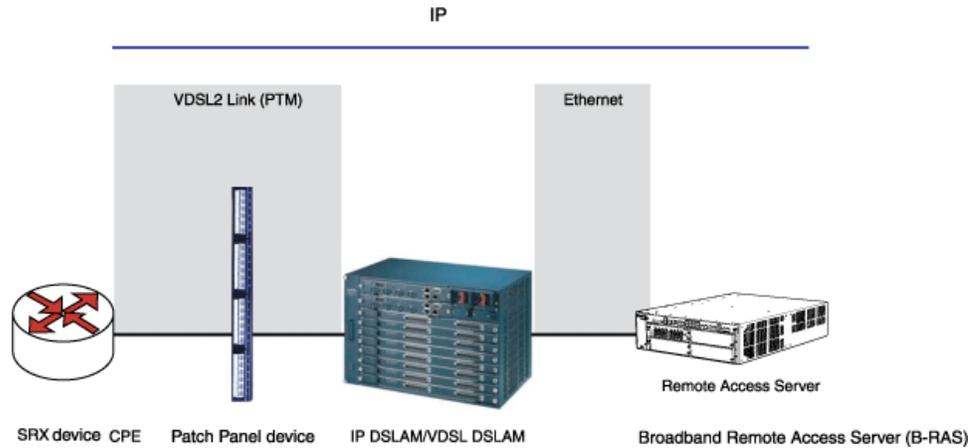
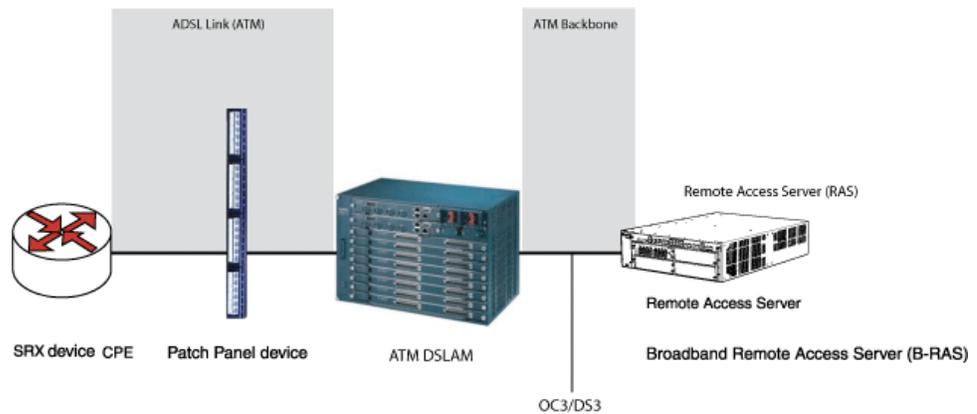


Figure 8 on page 77 shows a backward-compatible ADSL topology using ATM DSLAM.

Figure 8: Backward-Compatible ADSL Topology (ATM DSLAM)



Related Documentation

- [VDSL2 Interface Technology Overview on page 75](#)
- [VDSL2 Interface Support on SRX Series Devices on page 77](#)
- [Example: Configuring VDSL2 Interfaces \(Basic\) on page 81](#)
- [Example: Configuring VDSL2 Interfaces \(Detail\) on page 87](#)

VDSL2 Interface Support on SRX Series Devices

Supported Platforms [SRX320, SRX340](#)

The VDSL2 interface is supported on the SRX Series devices listed in [Table 14 on page 78](#).

Table 14: VDSL2 Annex A and Annex B Features

Features	POTS	ISDN
Devices	Integrated VDSL Module (SRX320-POTS) VDSL Mini-PIM (SRX320, SRX340)	Integrated VDSL Module (SRX320-ISDN)
Supported annex operating modes	Annex A and Annex B*	Annex B
Supported Bandplans	997/998	998
Supported standards	ITU-T G.993.2 and ITU-T G.993.5 (VDSL2)	ITU-T G.993.2 and ITU-T G.993.5 (VDSL2)
Used in	North American network implementations	European network implementations
ADSL backward compatibility	ADSL G992.5-A (ADSL Annex A)	ADSL G992.5-B (ADSL Annex B)

* Annex B support is not available on VDSL2 Mini-PIMs.

VDSL2 Interface Compatibility with ADSL Interfaces

VDSL2 interfaces on SRX Series devices are backward compatible with most ADSL interface standards. The VDSL2 interface uses Ethernet in the First Mile (EFM) mode or Packet Transfer Mode (PTM) and uses the named interface `pt-1/0/0`. In ADSL fallback mode, VDSL2 operates on the ATM encapsulation interface in the first mile and uses the named interface `at-1/0/0`.



NOTE:

- The VDSL2 interface has backward compatibility with ADSL/ADSL2/ADSL2+. The VDSL2 interface is represented by the `pt` interface when configured to function as VDSL2, and the ADSL interface is represented by the `at` interface when configured to function as ADSL.
- On VDSL2 interfaces, by default the `pt-1/0/0` interface is created when there is no configuration already created for either the `pt-1/0/0` or the `at-1/0/0` interface.



NOTE: It requires around 60 seconds to switch from VDSL2 to ADSL or from ADSL to VDSL2 operating modes.

Table 15 on page 79 lists VDSL2 operating modes and their backward compatibility with ADSL interface standards.

Table 15: VDSL2 Operating Mode Backward Compatibility with ADSL

VDSL2 Annex Type	Operating Modes	Description
VDSL2 Annex A interface (ADSL modes for Annex A only)	auto	Configures the ADSL interface to autonegotiate settings with the DSLAM located at the central office. For Annex A, the ADSL interface uses either ANSI T1.413 Issue II mode or ITU G.992.1 mode. NOTE: Automatic (auto) operating mode does not work when the DSLAM located at the central office is operating at ADSL2+ Annex M mode.
	ansi-dmt	Configures the ADSL interface to use ANSI T1.413 Issue II mode.
	itu-dmt	Configures the ADSL interface to use ITU G.992.1 mode.
	itu-dmt-bis	Configures the ADSL interface to use ITU G.992.3 mode. You can configure this mode only when it is supported on the DSLAM.
	adsl2plus	Configures the ADSL interface to use ITU G.992.5 mode. You can configure this mode only when it is supported on the DSLAM.
VDSL2 Annex B interface (ADSL modes for Annex B only)	auto	Configures the ADSL interface to autonegotiate settings with the DSLAM located at the central office. For Annex B, the ADSL interface trains in ITU G.992.1 mode.
	itu-dmt	Configures the ADSL interface to use ITU G.992.1 mode.
	itu-dmt-bis	Configures the ADSL interface to use ITU G.992.3 mode. You can configure this mode only when it is supported on the DSLAM.
	adsl2plus	Configures the ADSL interface to use ITU G.992.5 mode. You can configure this mode only when it is supported on the DSLAM.
	itu-annexb-ur2	Configures the ADSL line to use G.992.1 Deutsche Telekom UR-2 mode.

VDSL2 Interfaces Supported Profiles

A profile is a table that contains a list of preconfigured VDSL2 settings. [Table 16 on page 80](#) lists the different profiles supported on the VDSL2 interfaces and their properties.

Table 16: Supported Profiles on the VDSL2 Interfaces

Profiles	Data Rate
8a	50
8b	50
8c	50
8d	50
12a	68
12b	68
17a	100
Auto	Negotiated (based on operating mode)

VDSL2 Interfaces Supported Features

The following features are supported on the VDSL2 interfaces:

- ADSL/ADSL2/ADSL2+ backward compatibility with Annex A, Annex M support
- PTM or EFM (802.3ah) support
- Operation, Administration, and Maintenance (OAM) support for ADSL/ADSL2/ADSL2+ mode
- ATM quality of service (QoS) (supported only when the VDSL2 Mini-PIM is operating in ADSL2 mode)
- Multilink Point-to-Point Protocol (MLPPP) (supported only when the VDSL2 Mini-PIM is operating in ADSL2 mode)
- MTU size of 1514 bytes (maximum) in VDSL2 mode and 1496 bytes in ADSL mode.
- Support for maximum of 10 permanent virtual connections (PVCs) (only in ADSL/ADSL2/ADSL2+ mode)
- Dying gasp support (ADSL and VDSL2 mode)



NOTE: On SRX320 devices with VDSL2, ATM CoS VBR-related functionality cannot be tested.

Related Documentation

- [VDSL2 Interface Technology Overview on page 75](#)
- [VDSL2 Network Deployment Topology on page 76](#)
- [Example: Configuring VDSL2 Interfaces \(Basic\) on page 81](#)

- [Example: Configuring VDSL2 Interfaces \(Detail\) on page 87](#)

Example: Configuring VDSL2 Interfaces (Basic)

Supported Platforms [SRX320, SRX340](#)

This example shows how to configure the VDSL2 interfaces for SRX320, and SRX340 devices.

- [Requirements on page 81](#)
- [Overview on page 81](#)
- [Configuration on page 81](#)
- [Verifying the Configuration on page 83](#)

Requirements

Before you begin:

- Establish basic connectivity. See the *Quick Start Guide* for your device for factory default settings.
- Configure network interfaces as necessary. See “[Example: Creating an Ethernet Interface](#)” on page 122.

Overview

In this example, you create a VDSL2 interface called **pt-1/0/0**, specify the type of encapsulation, and set the VDSL2 profile to auto.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces pt-1/0/0 vdsl-options vdsl-profile auto
set interfaces pt-1/0/0 vlan-tagging
set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
set interfaces pt-1/0/0 unit 0 family inet dhcp
set interfaces pt-1/0/0 unit 0 vlan-id 100
```

Step-by-Step Procedure To configure the VDSL2 interfaces for the SRX320, and SRX340 devices and enable VLAN tagging:

1. Create an interface.


```
[edit]
user@host# edit interfaces pt-1/0/0
```
2. Set the type of VDSL2 profile.


```
[edit interfaces pt-1/0/0]
```

```
user@host# set vdsl-options vdsl-profile auto
```

3. Specify the logical unit to connect to this physical VDSL2 interface.

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

4. Specify the family protocol type.

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

5. To enable the DHCP client on the interface.

```
[edit interfaces pt-1/0/0]
user@host# set unit 0 family inet dhcp
```

6. Specify the type of encapsulation on the VDSL2 logical interface.

```
[edit interfaces pt-1/0/0]
user@host# set unit 0 encapsulation ppp-over-ether
```



NOTE: The VDSL2 interface supports PPPoE. You can also set no encapsulation for the VDSL2 interface.



NOTE: To configure VLAN tagging, continue the configuration with the next step.

7. To enable VLAN tagging on the pt interface.

```
[edit interfaces pt-1/0/0]
user@host# set interface pt-1/0/0 vlan-tagging
```

8. Specify the value of the VLAN ID to be configured.

```
[edit interfaces pt-1/0/0]
user@host# set interface pt-1/0/0 unit 0 vlan-id 100
```



NOTE: This feature is supported only on the pt interface, and the range of VLANs that can be configured is 0 to 4093.

Results From configuration mode, confirm your configuration by entering the **show interfaces pt-1/0/0** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces pt-1/0/0
vdsl-options {
vdsl-profile auto;
}
```

```

unit 0 {
  encapsulation ppp-over-ether;
  Family inet {
    address 100.100.100.1/24;
    dhcp;
  }
}

```



NOTE: When VLAN tagging is configured, the intended output is:

```

[edit]
user@host# show interfaces pt-1/0/0
vlan-tagging;
vdsl-options {
vdsl-profile auto;
}
unit 0 {
  encapsulation ppp-over-ether;
  vlan-id 100;
  Family inet {
    address 100.100.100.1/24;
    dhcp;
  }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verifying the Configuration

Confirm that the configuration is working properly.

- [Displaying the Configuration for VDSL2 Interface \(When Connected to the DSLAM Operating in Annex A Mode\) on page 83](#)
- [Displaying the Configuration for VDSL2 Interface \(When Connected to the DSLAM Operating in Annex B Mode\) on page 86](#)

Displaying the Configuration for VDSL2 Interface (When Connected to the DSLAM Operating in Annex A Mode)

Purpose Verify the command output.

Action From operational mode, enter the **show interfaces pt-1/0/0** command.

```

Physical interface: pt-1/0/0, Enabled, Physical link is Up
Interface index: 146, SNMP ifIndex: 524, Generation: 149
Type: PTM, Link-level type: Ethernet, MTU: 1496, VDSL mode, Speed: 45440kbps

Speed: VDSL2
Device flags : Present Running
Link flags : None
CoS queues : 8 supported, 8 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:b1:7e:85:84:ff
Last flapped : 2009-10-18 11:56:50 PDT (12:32:49 ago)
Statistics last cleared: 2009-10-19 00:29:37 PDT (00:00:02 ago)

```

```

Traffic statistics:
Input bytes : 22438962 97070256 bps
Output bytes : 10866024 43334088 bps
Input packets: 15141 8187 pps
Output packets: 7332 3655 pps
Input errors:
Errors: 0, Drops: 0, Policed discards: 0, L3 incompletes: 0,
L2 channel errors: 0, L2 mismatch timeouts: 0, Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, Aged packets: 0, MTU errors: 0,
Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets
0 best-effort 6759 6760 0
1 expedited-fo 0 0 0
2 assured-forw 0 0 0
3 network-cont 0 0 0
VDSL alarms : None
VDSL defects : None
VDSL media: Seconds Count State
LOF 0 0 OK
LOS 0 0 OK
LOM 0 0 OK
LOP 0 0 OK
LOCDI 0 0 OK
LOCDNI 0 0 OK
VDSL status:
Modem status : Showtime (Profile-17a)
VDSL profile : Profile-17a Annex A
Last fail code: None
Subfunction : 0x00
Seconds in showtime : 45171
VDSL Chipset Information: VTU-R VTU-C
Vendor Country : 0xb5 0xb5
Vendor ID : BDCM BDCM
Vendor Specific: 0x9385 0x9385
VDSL Statistics: VTU-R VTU-C
Attenuation (dB) : 0.0 0.0
Capacity used (%) : 0 0
Noise margin (dB) : 20.0 20.0
Output power (dBm) : 6.0 12.0
                                Interleave Fast Interleave Fast
Bit rate (kbps) :          100004          0 45440          0
CRC :                0          0 0          0
FEC :                0          0 0          0
HEC :                0          0 0          0
Packet Forwarding Engine configuration:
Destination slot: 0 (0x00)
CoS information:
Direction : Output
CoS transmit queue Bandwidth Buffer Priority
Limit
                                %          bps          %          usec
0 best-effort          95          43168000          95          0          low
none
3 network-control          5          2272000          5          0          low
none
Logical interface pt-1/0/0.0 (Index 71) (SNMP ifIndex 525) (Generation 136)
Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
Input bytes : 23789064
    
```

```

Output bytes : 10866024
Input packets: 16052
Output packets: 7332
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 23789064 97070256 bps
Output bytes : 10866024 43334088 bps
Input packets: 16052 8187 pps
Output packets: 7332 3655 pps
Security: Zone: Null
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: 1482, Generation: 169, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary

Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255,
Generation: 158

```

The output shows a summary of VDSL2 interface. Verify the following information:

- Status of interface pt-1/0/0 is displayed as Physical link is Up.
- Modem status is displayed as Showtime (Profile-17a).

- Time in seconds during which the interface stayed up is displayed as Seconds in showtime.
- Annex A indicates VDSL profile of the DSLAM connected at other end.

Displaying the Configuration for VDSL2 Interface (When Connected to the DSLAM Operating in Annex B Mode)

Purpose Verify the command output.

Action From operational mode, enter the **show interfaces pt-1/0/0** command.

```
Physical interface: pt-1/0/0, Enabled, Physical link is Up
  Interface index: 148, SNMP ifIndex: 536, Generation: 238
  Type: PTM, Link-level type: Ethernet, MTU: 1514, VDSL mode, Speed: 45439kbps
  Speed: VDSL2
  Device flags      : Present Running
  Link flags       : None
  CoS queues       : 8 supported, 8 maximum usable queues
  Hold-times       : Up 0 ms, Down 0 ms
  Current address: 00:1f:12:e4:df:20
  Last flapped    : 2011-05-13 07:34:33 PDT (00:46:33 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
  Input errors:
    Errors: 0, Drops: 0, Policed discards: 0, L3 incompletes: 0, L2 channel errors:
  0, L2 mismatch timeouts: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 3, Errors: 0, Drops: 0, Aged packets: 0, MTU errors: 0,
  Resource errors: 0
  VDSL alarms      : None
  VDSL defects     : None
  VDSL media:
    Seconds      Count  State
  LOF            177     0 OK
  LOS            177     0 OK
  LOM             0     0 OK
  LOP             0     0 OK
  LOCDI          0     0 OK
  LOCDNI         177     0 OK
  VDSL status:
  Modem status   : Showtime (Profile-17a)
  VDSL profile   : Auto Annex B
  Last fail code: None
  Subfunction    : 0x00
  Seconds in showtime: 2794
  VDSL Chipset Information:
    Vendor Country : 0xb5          0xb5
    Vendor ID      : BDCM          BDCM
    Vendor Specific: 0x9385        0x9395
  VDSL Statistics:
    Attenuation (dB) : 0.0          0.0
    Capacity used (%) : 0            0
    Noise margin (dB) : 18.5         9.5
    Output power (dBm) : 14.5         3.0
```

```

          Interleave      Fast Interleave      Fast
Bit rate (kbps) :          100015          0          45439          0
CRC              :              0          0              0          0
FEC              :              0          0              0          0
HEC              :              0          0              0          0
Packet Forwarding Engine configuration:
  Destination slot: 0 (0x00)
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority
Limit
  0 best-effort           95          43167050      95          0          low
none
  3 network-control       5          2271950       5          0          low
none

```

The output shows a summary of the VDSL2 interface. Verify the following information:

- Status of interface pt-1/0/0 is displayed as Physical link is Up.
- Modem status is displayed as Showtime (Profile-17a).
- Time in seconds during which the interface stayed up is displayed as Seconds in showtime.
- Annex B indicates the VDSL profile of the DSLAM connected at other end.

Related Documentation

- [Understanding Interfaces on page 3](#)
- [VDSL2 Interface Technology Overview on page 75](#)
- [Example: Configuring VDSL2 Interfaces \(Detail\) on page 87](#)
- [Example: Configuring VDSL2 Interfaces in ADSL Mode \(Detail\)](#)

Example: Configuring VDSL2 Interfaces (Detail)

Supported Platforms SRX320, SRX340

This example shows how to configure VDSL2 interfaces on SRX Series Services Gateways.

This example uses VDSL2 Mini-PIM installed on SRX320 devices. The information is also applicable to SRX320 devices (with VDSL2 Mini-PIMs).

- [Requirements on page 87](#)
- [Overview on page 88](#)
- [Configuration on page 89](#)
- [Verification on page 100](#)

Requirements

Before you begin:

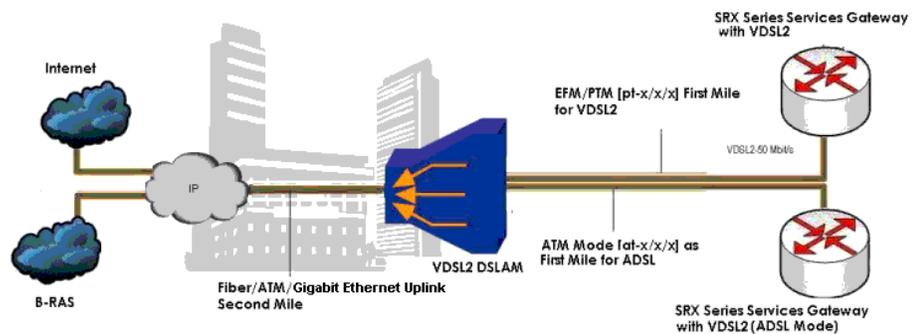
- Install Junos OS Release 10.1 or later on the SRX Series devices.
- Establish basic connectivity and set up and perform initial configuration. See the *Quick Start Guide* for your device for factory default settings.
- Install the VDSL2 Mini-PIM on the SRX320 device chassis.
- Connect the SRX320 device to a DSLAM.
- On VDSL2 Mini-PIMs, by default the **pt-1/0/0** interface is created when there is no configuration already created for either the **pt-1/0/0** or the **at-1/0/0** interface. You can switch to ADSL mode by just configuring **at-1/0/0**. If the configurations are already created for **pt-1/0/0** or **at-1/0/0**, then you need to deactivate **pt-1/0/0** before you create **at-1/0/0** or deactivate **at-1/0/0** to create **pt-1/0/0**.
- Make sure that you have deleted the previous configurations on **pt-1/0/0** and **pp0**.

Overview

This example uses SRX320 devices. The information is also applicable to SRX340 devices.

Figure 9 on page 88 shows typical SRX Series devices with VDSL2 Mini-PIM network connections.

Figure 9: SRX Series Device with VDSL2 Mini-PIMs in an End-to-End Deployment Scenario



In this example, you begin a new configuration on a VDSL2 Mini-PIM. You first deactivate previous interfaces and delete any old configuration from the device. Then you set the interfaces with the VDSL profile and the Layer 3 configuration for the end-to-end data path.

You then configure the PPPoE on the **pt-1/0/0** interface with a static IP address or CHAP authentication. You configure PPPoE on the **pt-1/0/0** interface with unnumbered IP address (PAP authentication or CHAP authentication).

Finally, you configure PPPoE on the **pt-1/0/0** interface with negotiated IP address (PAP authentication or CHAP authentication).

Configuration

- [Beginning a New Configuration on a VDSL2 Mini-PIM on page 89](#)
- [Configuring the VDSL2 Mini-PIM for End-to-End Data Path on page 90](#)
- [Configuring PPPoE on the pt-1/0/0 Interface with a Static IP Address on page 91](#)
- [Configuring PPPoE on the pt-1/0/0 Interface with a Static IP Address \(CHAP Authentication\) on page 93](#)
- [Configuring PPPoE on the pt-x/x/x Interface with Unnumbered IP \(PAP Authentication\) on page 94](#)
- [Configuring PPPoE on the pt-1/0/0 Interface with Unnumbered IP \(CHAP Authentication\) on page 96](#)
- [Configuring PPPoE on the pt-1/0/0 Interface with Negotiated IP \(PAP Authentication\) on page 97](#)
- [Configuring PPPoE on the pt-1/0/0 Interface with Negotiated IP \(CHAP Authentication\) on page 99](#)

Beginning a New Configuration on a VDSL2 Mini-PIM

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
deactivate interface pt-1/0/0
deactivate interface at-1/0/0
delete interface pt-1/0/0
delete interface pp0
```

Step-by-Step Procedure

To begin a new configuration on a VDSL2 Mini-PIM:

1. Deactivate any previous interfaces.

```
[edit]
user@host# deactivate interface pt-1/0/0
user@host# deactivate interface at-1/0/0
```

2. Delete any old configurations.

```
[edit]
user@host# delete interface pt-1/0/0
user@host# delete interface pp0
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

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

```
user@host# run show chassis fpc
Temp CPU Utilization (%) Memory Utilization
(%)
Slot State (C) Total Interrupt DRAM (MB) Heap Buffer
0 Online ----- CPU less FPC -----
1 Online ----- CPU less FPC -----
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the VDSL2 Mini-PIM for End-to-End Data Path

CLI Quick Configuration To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
set interfaces pt-1/0/0 unit 0 family inet address 11.11.11.1/24
```

Step-by-Step Procedure To configure the VDSL2 Mini-PIM for end-to-end data path:

1. Configure the interfaces with the VDSL profile and the Layer 3 configuration for end-to-end data path.

```
[edit]
user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
user@host# set interfaces pt-1/0/0 unit 0 family inet address 11.11.11.1/24
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces pt-1/0/0** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces pt-1/0/0
vdsl-options {
  vdsl-profile 17a;
}
unit 0 {
  family inet {
    address 11.11.11.1/24;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring PPPoE on the pt-1/0/0 Interface with a Static IP Address

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
user@host# set interfaces pp0 unit 0 ppp-options pap access-profile pap_prof local-name
  locky local-password india passive
user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
  auto-reconnect 120 client
user@host# set interfaces pp0 unit 0 family inet address 10.1.1.6/24
user@host# set access profile pap_prof authentication-order password client cuttack
  pap-password india
```



NOTE: To configure VLAN tagging while configuring PPPoE on the pt-1/0/0 interface with

- Static IP address
- Static IP address (CHAP authentication)
- Unnumbered IP address (PAP Authentication)
- Unnumbered IP address (CHAP Authentication)
- Negotiated IP address (PAP Authentication)
- Negotiated IP address (CHAP Authentication)

the following commands must be included at **[edit]** hierarchy level:

```
set interfaces pt-1/0/0 vlan-tagging
set interfaces pt-1/0/0 unit 0 vlan-id 100
```

Step-by-Step Procedure

To configure the PPPoE on the pt-1/0/0 interface with a static IP address:

1. Configure the VDSL options and encapsulation for the interface.

```
[edit]
user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
```

2. Configure the PPP options for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 ppp-options pap access-profile pap_prof
user@host# set interfaces pp0 unit 0 ppp-options pap local-name locky
user@host# set interfaces pp0 unit 0 ppp-options pap local-password india
user@host# set interfaces pp0 unit 0 ppp-options pap passive
```

3. Configure the PPPoE options for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
user@host# set interfaces pp0 unit 0 pppoe-options auto-reconnect 120
user@host# set interfaces pp0 unit 0 pppoe-options client
```

4. Configure the IP address for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 family inet address 10.1.1.6/24
```

5. Configure the access profile for the interface.

```
[edit]
user@host# set access profile pap_prof authentication-order password
user@host# set access profile pap_prof client cuttack pap-password india
```

Results From configuration mode, confirm your configuration by entering the **show interfaces pp0**, **show interfaces pt-1/0/0** and **show access profile pap_prof** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces pp0
unit 0 {
  ppp-options {
    pap {
      access-profile pap_prof;
      local-name locky;
      local-password "$ABC123"; ## SECRET-DATA
    }
    passive;
  }
  pppoe-options {
    underlying-interface pt-1/0/0.0;
    auto-reconnect 120;
    client;
  }
  family inet {
    address 10.1.1.6/24;
  }
}
[edit]
user@host# show interfaces pt-1/0/0
vdsl-options {
  vdsl-profile 17a;
}
unit 0 {
  encapsulation ppp-over-ether;
}
[edit]
user@host# show access profile pap_prof
authentication-order password;
client cuttack pap-password "$ABC123"; ## SECRET-DATA
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring PPPoE on the pt-1/0/0 Interface with a Static IP Address (CHAP Authentication)

CLI Quick Configuration To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
user@host# set interfaces pp0 unit 0 ppp-options chap default-chap-secret india
local-name locky passive
user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
auto-reconnect 120 client
user@host# set interfaces pp0 unit 0 family inet address 10.1.1.6/24
```

Step-by-Step Procedure To configure the PPPoE on the pt-1/0/0 interface with a static IP address (CHAP authentication):

1. Configure the VDSL options and encapsulation for the interface.

```
[edit]
user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
```

2. Configure the PPP options for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 ppp-options chap default-chap-secret india
user@host# set interfaces pp0 unit 0 ppp-options chap local-name locky
user@host# set interfaces pp0 unit 0 ppp-options chap passive
```

3. Configure the PPPoE options for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
user@host# set interfaces pp0 unit 0 pppoe-options auto-reconnect 120
user@host# set interfaces pp0 unit 0 pppoe-options client
```

4. Configure the IP address for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 family inet address 10.1.1.6/24
```

Results From configuration mode, confirm your configuration by entering the **show interfaces pt-1/0/0** and **show interfaces pp0** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces pt-1/0/0
vdsl-options {
  vdsl-profile 17a;
}
unit 0 {
  encapsulation ppp-over-ether;
}
```

```
[edit]
user@host# show interfaces pp0
unit 0 {
  ppp-options {
    chap {
      default-chap-secret "$ABC123"; ## SECRET-DATA
    }
    local-name locky;
    passive;
  }
  pppoe-options {
    underlying-interface pt-1/0/0.0;
    auto-reconnect 120;
    client;
  }
  family inet {
    address 10.1.1.6/24;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring PPPoE on the pt-x/x/x Interface with Unnumbered IP (PAP Authentication)

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
user@host# set interfaces lo0 unit 0 family inet address 10.1.1.24/32
user@host# set interfaces pp0 unit 0 ppp-options pap access-profile pap_prof local-name locky local-password india passive
user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0 auto-reconnect 120 client
user@host# set interfaces pp0 unit 0 family inet unnumbered-address lo0.0 destination 10.1.1.1
user@host# set access profile pap_prof authentication-order password client cuttack pap-password india
```

Step-by-Step Procedure

To configure PPPoE on the pt-1/0/0 interface with unnumbered IP (PAP authentication):

1. Configure the VDSL options and encapsulation for the interface.

```
[edit]
user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
```

2. Configure the IP address for the interface.

```
[edit]
user@host# set interfaces lo0 unit 0 family inet address 10.1.1.24/32
```

3. Configure the PPP options for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 ppp-options pap access-profile pap_prof
user@host# set interfaces pp0 unit 0 ppp-options pap local-name locky
user@host# set interfaces pp0 unit 0 ppp-options pap local-password india
user@host# set interfaces pp0 unit 0 ppp-options pap passive
```

4. Configure the PPPoE options for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
user@host# set interfaces pp0 unit 0 pppoe-options auto-reconnect 120
user@host# set interfaces pp0 unit 0 pppoe-options client
```

5. Configure the unnumbered address and destination for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 family inet unnumbered-address lo0.0
user@host# set interfaces pp0 unit 0 family inet unnumbered-address destination
10.1.1.1
```

6. Configure the access profile for the interface.

```
[edit]
user@host# set access profile pap_prof authentication-order password
user@host# set access profile pap_prof client cuttack pap-password india
```

Results From configuration mode, confirm your configuration by entering the **show interfaces lo0**, **show interfaces pt-1/0/0**, and **show interfaces pp0** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces lo0
unit 0 {
family inet {
address 10.1.1.24/32;
}
}
[edit]
user@host# show interfaces pt-1/0/0
vdsl-options {
vdsl-profile 17a;
}
unit 0 {
encapsulation ppp-over-ether;
}
[edit]
user@host# show interfaces pp0
unit 0 {
ppp-options {
pap {
access-profile pap_prof;
local-name locky;
local-password "$ABC123"; ## SECRET-DATA
passive;
}
}
}
```

```

pppoe-options {
  underlying-interface pt-1/0/0.0;
  auto-reconnect 120;
  client;
}
family inet {
  unnumbered-address lo0.0 destination 10.1.1.1;
}
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring PPPoE on the pt-1/0/0 Interface with Unnumbered IP (CHAP Authentication)

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```

user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
user@host# set interfaces lo0 unit 0 family inet address 10.1.1.24/32
user@host# set interfaces pp0 unit 0 ppp-options chap default-chap-secret india
  local-name locky passive
user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
  auto-reconnect 120 client
user@host# set interfaces pp0 unit 0 family inet unnumbered-address lo0.0 destination
10.1.1.1

```

Step-by-Step Procedure

To configure PPPoE on the pt-1/0/0 interface with unnumbered IP (CHAP authentication):

1. Configure the VDSL options and encapsulation for the interface.


```

[edit]
user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether

```
2. Configure the IP address for the interface.


```

[edit]
user@host# set interfaces lo0 unit 0 family inet address 10.1.1.24/32

```
3. Configure the PPP options for the interface.


```

[edit]
user@host# set interfaces pp0 unit 0 ppp-options chap default-chap-secret india
user@host# set interfaces pp0 unit 0 ppp-options chap local-name locky
user@host# set interfaces pp0 unit 0 ppp-options chap passive

```
4. Configure the PPPoE options for the interface.


```

[edit]
user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
user@host# set interfaces pp0 unit 0 pppoe-options auto-reconnect 120
user@host# set interfaces pp0 unit 0 pppoe-options client

```
5. Configure the unnumbered address and destination for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 family inet unnumbered-address lo0.0
user@host# set interfaces pp0 unit 0 family inet unnumbered-address destination
10.1.1.1
```

Results From configuration mode, confirm your configuration by entering the **show interfaces pp0**, **show interfaces pt-1/0/0**, and **show interfaces lo0** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces pp0
unit 0 {
  ppp-options {
    chap {
      default-chap-secret "$ABC123"; ## SECRET-DATA
      local-name locky;
      passive;
    }
  }
  pppoe-options {
    underlying-interface pt-1/0/0.0;
    auto-reconnect 120;
    client;
  }
  family inet {
    unnumbered-address lo0.0 destination 10.1.1.1;
  }
}
[edit]
user@host# show interfaces pt-1/0/0
vdsl-options {
  vdsl-profile 17a;
}
unit 0 {
  encapsulation ppp-over-ether;
}
[edit]
user@host# show interfaces lo0
unit 0 {
  family inet {
    address 10.1.1.24/32;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring PPPoE on the pt-1/0/0 Interface with Negotiated IP (PAP Authentication)

CLI Quick Configuration To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```

user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
user@host# set interfaces pp0 unit 0 ppp-options pap access-profile my_prf local-name
purple local-password <password> passive
user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
auto-reconnect 120 client
user@host# set interfaces pp0 unit 0 family inet negotiate-address
user@host# set access profile my_prf authentication-order password
user@host# set access profile my_prf

```

Step-by-Step Procedure

To configure PPPoE on the pt-1/0/0 interface with negotiated IP (PAP authentication):

1. Configure the VDSL options and encapsulation for the interface.


```

[edit]
user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether

```
2. Configure the PPP options for the interface.


```

[edit]
user@host# set interfaces pp0 unit 0 ppp-options pap access-profile my_prf
user@host# set interfaces pp0 unit 0 ppp-options pap local-name purple
user@host# set interfaces pp0 unit 0 ppp-options pap local-password <password>
user@host# set interfaces pp0 unit 0 ppp-options pap passive

```
3. Configure the PPPoE options for the interface.


```

[edit]
user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
user@host# set interfaces pp0 unit 0 pppoe-options auto-reconnect 120
user@host# set interfaces pp0 unit 0 pppoe-options client

```
4. Configure the negotiated IP address for the interface.


```

[edit]
user@host# set interfaces pp0 unit 0 family inet negotiate-address

```
5. Configure the access profile for the interface.


```

[edit]
user@host# set access profile my_prf authentication-order password
user@host# set access profile my_prf

```

Results

From configuration mode, confirm your configuration by entering the **show interfaces pt-1/0/0**, **show interfaces pp0**, and **show access profile my_prf** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```

[edit]
user@host# show interfaces pt-1/0/0
vdsl-options {
vdsl-profile 17a;
}
unit 0 {
encapsulation ppp-over-ether;
}
[edit]

```

```

user@host# show interfaces pp0
unit 0 {
  ppp-options {
    pap {
      access-profile my_prf;
      local-name purple;
      local-password "$ABC123"; ## SECRET-DATA
    }
    passive;
  }
  pppoe-options {
    underlying-interface pt-1/0/0.0;
    auto-reconnect 120;
  }
  client;
}
  family inet {
    negotiate-address;
  }
}
[edit]
user@host# show access profile my_prf
authentication-order password;

```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring PPPoE on the pt-1/0/0 Interface with Negotiated IP (CHAP Authentication)

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```

user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
user@host# set interfaces pp0 unit 0 ppp-options chap default-chap-secret <password>
  local-name purple passive
user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
  auto-reconnect 120 client
user@host# set interfaces pp0 unit 0 family inet negotiate-address

```

Step-by-Step Procedure

To configure PPPoE on the pt-1/0/0 interface with negotiated IP (CHAP authentication):

1. Configure the VDSL options and encapsulation for the interface.

```

[edit]
user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether

```

2. Configure the PPP options for the interface.

```

[edit]
user@host# set interfaces pp0 unit 0 ppp-options chap default-chap-secret
  <password>
user@host# set interfaces pp0 unit 0 ppp-options chap local-name purple
user@host# set interfaces pp0 unit 0 ppp-options chap passive

```

3. Configure the PPPoE options for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
user@host# set interfaces pp0 unit 0 pppoe-options auto-reconnect 120
user@host# set interfaces pp0 unit 0 pppoe-options client
```

4. Configure the negotiated IP address for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 family inet negotiate-address
```

Results From configuration mode, confirm your configuration by entering the **show interfaces pp0** and **show interfaces pt-1/0/0** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces pp0
unit 0 {
  ppp-options {
    chap {
      default-chap-secret "$ABC123"; ## SECRET-DATA
      local-name purple;
      passive;
    }
  }
  pppoe-options {
    underlying-interface pt-1/0/0.0;
    auto-reconnect 120;
    client;
  }
  family inet {
    negotiate-address;
  }
}
[edit]
user@host# show interfaces pt-1/0/0
vdsl-options {
  vdsl-profile 17a;
}
unit 0 {
  encapsulation ppp-over-ether;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying the Configuration on page 101](#)
- [Verifying the VDSL2 Mini-PIM for End-to-End Data Path on page 103](#)
- [Verifying PPPoE on the pt-1/0/0 Interface with a Static IP Address on page 106](#)

- Verifying PPPoE on the pt-1/0/0 Interface with a Static IP Address (CHAP Authentication) on page 107
- Verifying PPPoE on the pt-1/0/0 Interface with Unnumbered IP (PAP Authentication) on page 108
- Verifying PPPoE on the pt-1/0/0 Interface with Unnumbered IP (CHAP Authentication) on page 109
- Verifying PPPoE on the pt-1/0/0 Interface with Negotiated IP (PAP Authentication) on page 110
- Verifying PPPoE on the pt-1/0/0 Interface with Negotiated IP (CHAP Authentication) on page 111

Verifying the Configuration

Purpose Verify the FPC status and the command output.

- Action**
1. Verify the FPC status by entering the **show chassis fpc** command. The output should display FPC status as online.

```
user@host# run show chassis fpc
Temp CPU Utilization (%) Memory Utilization
(%)
Slot State (C) Total Interrupt DRAM (MB) Heap Buffer
0 Online ----- CPU less FPC -----
1 Online ----- CPU less FPC -----
```



NOTE: The VDSL2 Mini-PIM is installed in the first slot of the SRX320 device chassis; therefore, the FPC used here is fpc 1. For SRX340 devices, the FPC used will be fpc 1, fpc 2, fpc 3, or fpc 4.

2. Enter **run show interface pt-1/0/0** and verify the following information in the command output:
 - Status of interface pt-1/0/0 is displayed as physical link is up.
 - Modem status is displayed as Showtime (Profile-17a).
 - Time in seconds during which the interface stayed up is displayed as Seconds in Showtime.
 - VDSL profile of DSLAM is displayed as Auto Annex A.

```
Physical interface: pt-1/0/0, Enabled, Physical link is Up
Interface index: 146, SNMP ifIndex: 524, Generation: 149
Type: PTM, Link-level type: Ethernet, MTU: 1496, VDSL mode, Speed: 45440kbps

Speed: VDSL2
Device flags : Present Running
Link flags : None
CoS queues : 8 supported, 8 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:b1:7e:85:84:ff
Last flapped : 2009-10-18 11:56:50 PDT (12:32:49 ago)
```

```

Statistics last cleared: 2009-10-19 00:29:37 PDT (00:00:02 ago)
Traffic statistics:
  Input bytes : 22438962 97070256 bps
  Output bytes : 10866024 43334088 bps
  Input packets: 15141 8187 pps
  Output packets: 7332 3655 pps
Input errors:
  Errors: 0, Drops: 0, Policed discards: 0, L3 incompletes: 0,
  L2 channel errors: 0, L2 mismatch timeouts: 0, Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Aged packets: 0, MTU errors:
0,
  Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets
  0 best-effort 6759 6760 0
  1 expedited-fo 0 0 0
  2 assured-forw 0 0 0
  3 network-cont 0 0 0
VDSL alarms : None
VDSL defects : None
VDSL media: Seconds Count State
  LOF 0 0 OK
  LOS 0 0 OK
  LOM 0 0 OK
  LOP 0 0 OK
  LOCDI 0 0 OK
  LOCDNI 0 0 OK
VDSL status:
  Modem status : Showtime (Profile-17a)
  VDSL profile : Profile-17a Annex A
  Last fail code: None
  Subfunction : 0x00
  Seconds in showtime : 45171
VDSL Chipset Information: VTU-R VTU-C
  Vendor Country : 0xb5 0xb5
  Vendor ID : BDCM BDCM
  Vendor Specific: 0x9385 0x9385
VDSL Statistics: VTU-R VTU-C
  Attenuation (dB) : 0.0 0.0
  Capacity used (%) : 0 0
  Noise margin (dB) : 20.0 20.0
  Output power (dBm) : 6.0 12.0
          Interleave Fast Interleave Fast
  Bit rate (kbps) :    100004    0    45440    0
  CRC :              0          0    0          0
  FEC :              0          0    0          0
  HEC :              0          0    0          0
Packet Forwarding Engine configuration:
  Destination slot: 0 (0x00)
CoS information:
  Direction : Output
  CoS transmit queue Bandwidth Buffer Priority
Limit
          %          bps          %          usec
  0 best-effort      95      43168000  95          0      low
none
  3 network-control  5        2272000   5          0      low
none
Logical interface pt-1/0/0.0 (Index 71) (SNMP ifIndex 525) (Generation 136)
  Flags: SNMP-Traps Encapsulation: ENET2

```

```

Traffic statistics:
  Input bytes : 23789064
  Output bytes : 10866024
  Input packets: 16052
  Output packets: 7332
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 23789064 97070256 bps
  Output bytes : 10866024 43334088 bps
  Input packets: 16052 8187 pps
  Output packets: 7332 3655 pps
Security: Zone: Null
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: 1482, Generation: 169, Route table: 0
  Flags: None
  Addresses, Flags: Is-Preferred Is-Primary

Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255,
Generation: 158

```

Verifying the VDSL2 Mini-PIM for End-to-End Data Path

Purpose Verify the interface status and check traffic statistics.

- Action** 1. Verify interface status by using the **show interface terse** command and test end-to-end data path connectivity by sending the ping packets to the remote end IP address.

```

user@host# run show interfaces pt-1/0/0 terse
Interface          Admin Link Proto  Local          Remote
pt-1/0/0           up    up
pt-1/0/0.0         up    up    inet    11.11.11.1/24

[edit]
user@host# run ping 11.11.11.2 count 1000 rapid
PING 11.11.11.2 (11.11.11.2): 56 data bytes
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
- 11.11.11.2 ping statistics ---
1000 packets transmitted, 1000 packets received, 0% packet loss
round-trip min/avg/max/stddev = 16.109/17.711/28.591/2.026 ms
    
```

2. Verify the VDSL2 interface configuration and check the traffic statistics.

```

user@host# run show interfaces pt-1/0/0 extensive
Physical interface: pt-1/0/0, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 524, Generation: 197
  Type: PTM, Link-level type: Ethernet, MTU: 1496, VDSL mode, Speed: 45440kbps

Speed: VDSL2
Device flags   : Present Running
Link flags     : None
CoS queues    : 8 supported, 8 maximum usable queues
Hold-times    : Up 0 ms, Down 0 ms
Current address: 00:b1:7e:85:84:ff
Last flapped  : 2009-10-28 00:36:29 PDT (00:12:03 ago)
Statistics last cleared: 2009-10-28 00:47:56 PDT (00:00:36 ago)
Traffic statistics:
  Input bytes   :                84000          0 bps
  Output bytes  :               138000          0 bps
  Input packets:                 1000          0 pps
  Output packets:                1000          0 pps
Input errors:
  Errors: 0, Drops: 0, Policed discards: 0, L3 incompletes: 0, L2 channel
errors: 0, L2 mismatch timeouts: 0, Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Aged packets: 0, MTU errors:
0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort                1000                1000
0
  1 expedited-fo                  0                    0
0
  2 assured-forw                  0                    0
0
  3 network-cont                  0                    0
0
VDSL alarms   : None
VDSL defects  : None
VDSL media:
  Seconds      Count  State
LOF             0      0 OK
LOS             0      0 OK
LOM             0      0 OK
LOP             0      0 OK
LOCDI           0      0 OK
LOCDNI          0      0 OK
    
```

VDSL status:

Modem status : Showtime (Profile-17a)
 VDSL profile : Profile-17a Annex A
 Last fail code: None
 Subfunction : 0x00
 Seconds in showtime : 723

VDSL Chipset Information:	VTU-R	VTU-C
Vendor Country :	0xb5	0xb5
Vendor ID :	BDCM	BDCM
Vendor Specific:	0x9385	0x9385
VDSL Statistics:	VTU-R	VTU-C
Attenuation (dB) :	0.0	0.0
Capacity used (%) :	0	0
Noise margin (dB) :	16.0	20.0
Output power (dBm) :	5.0	13.0

	Interleave	Fast	Interleave	Fast
0 Bit rate (kbps) :	100004	0	45440	
0 CRC :	0	0	0	
0 FEC :	0	0	0	
0 HEC :	0	0	0	

Packet Forwarding Engine configuration:

Destination slot: 0 (0x00)

CoS information:

Direction :	CoS transmit queue	Bandwidth	Buffer	Priority	
Limit	%	bps	%	usec	
0 best-effort	95	43168000	95	0	low
none					
3 network-control	5	2272000	5	0	low
none					

Logical interface pt-1/0/0.0 (Index 72) (SNMP ifIndex 521) (Generation 158)

Flags: SNMP-Traps Encapsulation: ENET2

Traffic statistics:

Input bytes : 84000
 Output bytes : 98000
 Input packets: 1000
 Output packets: 1000

Local statistics:

Input bytes : 84000
 Output bytes : 98000
 Input packets: 1000
 Output packets: 1000

Transit statistics:

Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0

Security: Zone: Null

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: 1482, Generation: 169, Route table: 0
  Flags: None
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 11.11.11/24, Local: 11.11.11.1, Broadcast: 11.11.11.255,
  Generation: 189

```

Verifying PPPoE on the pt-1/0/0 Interface with a Static IP Address

Purpose Verify the interface output and the end-to-end data path.

Action 1. Verify the interface output.

```

user@host# run show interfaces pp0
Physical interface: pp0, Enabled, Physical link is Up
Interface index: 128, SNMP ifIndex: 510
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 packets : 0
Output packets: 0

Logical interface pp0.0 (Index 71) (SNMP ifIndex 522)
Flags: Hardware-Down Point-To-Point SNMP-Traps 0x0 Encapsulation: PPPoE
PPPoE:
  State: SessionDown, Session ID: None,
  Configured AC name: None, Service name: None,
  Auto-reconnect timeout: 120 seconds, Idle timeout: Never,
  Underlying interface: pt-1/0/0.0 (Index 69)
Input packets : 57

```

```

Output packets: 56
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive: Input: 22 (00:00:40 ago), Output: 25 (00:00:04 ago)
LCP state: Down
NCP state: inet: Down, inet6: Not-configured, iso: Not-configured, mp1s:
Not-configured
CHAP state: Closed
PAP state: Closed
Security: Zone: Null
Protocol inet, MTU: 1492
Flags: Protocol-Down
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
Destination: 10.1.1/24, Local: 10.1.1.6

```

2. Verify the end-to-end data path on the interface.

```

user@host# run show interfaces pt-1/0/0 terse
Interface      Admin Link Proto  Local      Remote
pt-1/0/0       up    up
pt-1/0/0.0     up    up

[edit]
user@host# run show interfaces pp0 terse
Interface      Admin Link Proto  Local      Remote
pp0            up    up
pp0.0          up    up  inet   10.1.1.6/24

[edit]
user@host# run ping 10.1.1.1 count 100 rapid
PING 10.1.1.1 (10.1.1.1): 56 data bytes
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
- 10.1.1.1 ping statistics ---
100 packets transmitted, 100 packets received, 0% packet loss
round-trip min/avg/max/stddev = 14.669/15.649/21.655/1.740 ms

```

Verifying PPPoE on the pt-1/0/0 Interface with a Static IP Address (CHAP Authentication)

Purpose Verify the interface status and check the end-to-end data path connectivity.

- Action** 1. Verify the interface status.

```

user@host# run show interfaces pp0
Physical interface: pp0, Enabled, Physical link is Up
Interface index: 128, SNMP ifIndex: 510
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 packets  : 0
Output packets: 0

Logical interface pp0.0 (Index 70) (SNMP ifIndex 522)
Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: PPPoE
PPPoE:
State: SessionUp, Session ID: 31,
Session AC name: cuttack, Remote MAC address: 00:03:6c:c8:8c:55,
Configured AC name: None, Service name: None,

```

```

        Auto-reconnect timeout: 120 seconds, Idle timeout: Never,
        Underlying interface: pt-1/0/0.0 (Index 69)
    Input packets : 12
    Output packets: 10
    Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
    Keepalive: Input: 1 (00:00:08 ago), Output: 0 (never)
    LCP state: Opened
    NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mp1s:
    Not-configured
    CHAP state: Success
    PAP state: Closed
    Security: Zone: Null
    Protocol inet, MTU: 1492
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.1.1/24, Local: 10.1.1.6
    
```

2. Verify the interface and check the end-to-end data path connectivity.

```

user@host# run show interfaces pt-1/0/0 terse
Interface      Admin Link Proto  Local      Remote
pt-1/0/0       up      up
pt-1/0/0.0     up      up

[edit]
user@host# run show interfaces pp0 terse
Interface      Admin Link Proto  Local      Remote
pp0            up      up
pp0.0          up      up   inet    10.1.1.6/24

[edit]
user@host# run ping 10.1.1.1 count 100 rapid
PING 10.1.1.1 (10.1.1.1): 56 data bytes

!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
--- 10.1.1.1 ping statistics ---
100 packets transmitted, 100 packets received, 0% packet loss
round-trip min/avg/max/stddev = 14.608/15.466/25.939/1.779 ms
    
```

Verifying PPPoE on the pt-1/0/0 Interface with Unnumbered IP (PAP Authentication)

Purpose Verify the interface status and the end-to-end data path testing.

- Action** 1. Verify the interface status.

```

user@host# run show interfaces pp0
Physical interface: pp0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 510
  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 packets : 0
  Output packets: 0

Logical interface pp0.0 (Index 72) (SNMP ifIndex 522)
  Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: PPPoE
  PPPoE:
    
```

```

State: SessionUp, Session ID: 33,
Session AC name: cuttack, Remote MAC address: 00:03:6c:c8:8c:55,
Configured AC name: None, Service name: None,
Auto-reconnect timeout: 120 seconds, Idle timeout: Never,
Underlying interface: pt-1/0/0.0 (Index 69)
Input packets : 22
Output packets: 20
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive: Input: 1 (00:00:08 ago), Output: 0 (never)
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mp1s:
Not-configured
CHAP state: Closed
PAP state: Success
Security: Zone: Null
Protocol inet, MTU: 1492
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.1.1.1, Local: 10.1.1.24

```

2. Verify the end-to-end data path testing.

```

user@host# run show interfaces pt-1/0/0 terse
Interface      Admin Link Proto  Local          Remote
pt-1/0/0       up    up
pt-1/0/0.0     up    up

[edit]
user@host# run show interfaces pp0 terse
Interface      Admin Link Proto  Local          Remote
pp0            up    up
pp0.0          up    up  inet  10.1.1.24      --> 10.1.1.1

[edit]
user@host# run ping 10.1.1.1 count 100 rapid
PING 10.1.1.1 (10.1.1.1): 56 data bytes
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
--- 10.1.1.1 ping statistics ---
100 packets transmitted, 100 packets received, 0% packet loss
round-trip min/avg/max/stddev = 14.584/15.503/21.204/1.528 ms

```

Verifying PPPoE on the pt-1/0/0 Interface with Unnumbered IP (CHAP Authentication)

Purpose Verify the interface status and end-to-end data path testing on the PPPoE interface.

Action 1. Verify the interface status.

```

user@host# run show interfaces pp0
Physical interface: pp0, Enabled, Physical link is Up
Interface index: 128, SNMP ifIndex: 510
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 packets  : 0
Output packets: 0

Logical interface pp0.0 (Index 70) (SNMP ifIndex 522)

```

```

Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: PPPoE
PPPoE:
  State: SessionUp, Session ID: 35,
  Session AC name: cuttack, Remote MAC address: 00:03:6c:c8:8c:55,
  Configured AC name: None, Service name: None,
  Auto-reconnect timeout: 120 seconds, Idle timeout: Never,
  Underlying interface: pt-1/0/0.0 (Index 69)
Input packets : 25
Output packets: 22
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive: Input: 2 (00:00:10 ago), Output: 2 (00:00:02 ago)
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mp1s:
Not-configured
CHAP state: Success
PAP state: Closed
Security: Zone: Null
Protocol inet, MTU: 1492
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.1.1.1, Local: 10.1.1.24
    
```

2. Verify the end-to-end data path testing on the PPPoE interface.

```

user@host# run show interfaces pt-1/0/0 terse
Interface      Admin Link Proto  Local      Remote
pt-1/0/0       up    up
pt-1/0/0.0     up    up

[edit]
user@host# run show interfaces pp0 terse
Interface      Admin Link Proto  Local      Remote
pp0            up    up
pp0.0          up    up  inet   10.1.1.24  --> 10.1.1.1

[edit]
user@host# run ping 10.1.1.1 count 100 rapid
PING 10.1.1.1 (10.1.1.1): 56 data bytes
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
-- 10.1.1.1 ping statistics ---
100 packets transmitted, 100 packets received, 0% packet loss
round-trip min/avg/max/stddev = 14.585/16.025/22.354/2.019 ms
    
```

Verifying PPPoE on the pt-1/0/0 Interface with Negotiated IP (PAP Authentication)

Purpose Verify the PPPoE interface status and the end-to-end data path connectivity.

- Action** 1. Verify the PPPoE interface status.

```

user@host# run show interfaces pp0
Physical interface: pp0, Enabled, Physical link is Up
Interface index: 128, SNMP ifIndex: 510
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 packets  : 0
Output packets: 0
    
```

```

Logical interface pp0.0 (Index 72) (SNMP ifIndex 522)
Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: PPPoE
PPPoE:
  State: SessionUp, Session ID: 4,
  Session AC name: belur, Remote MAC address: 00:90:1a:43:18:d1,
  Configured AC name: None, Service name: None,
  Auto-reconnect timeout: 120 seconds, Idle timeout: Never,
  Underlying interface: pt-1/0/0.0 (Index 69)
  Input packets : 18
  Output packets: 18
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 0 (never), Output: 11 (00:00:01 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
  CHAP state: Closed
  PAP state: Success
  Security: Zone: Null
  Protocol inet, MTU: 1474
  Flags: Negotiate-Address
  Addresses, Flags: Kernel Is-Preferred Is-Primary
  Destination: 12.12.12.1, Local: 12.12.12.11

```

2. Verify the end-to-end data path connectivity.

```

user@host# run show interfaces pt-1/0/0 terse
Interface      Admin Link Proto  Local      Remote
pt-1/0/0       up    up
pt-1/0/0.0     up    up

[edit]
user@host# run show interfaces pp0 terse
Interface      Admin Link Proto  Local      Remote
pp0            up    up
pp0.0          up    up  inet    12.12.12.11  --> 12.12.12.1

[edit]
user@host# run ping 12.12.12.1 count 100 rapid
PING 12.12.12.1 (12.12.12.1): 56 data bytes
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
--- 12.12.12.1 ping statistics ---
100 packets transmitted, 100 packets received, 0% packet loss
round-trip min/avg/max/stddev = 16.223/17.692/24.359/2.292 ms

```

Verifying PPPoE on the pt-1/0/0 Interface with Negotiated IP (CHAP Authentication)

Purpose Verify the interface status and the end-to-end data path connectivity.

Action 1. Verifying the interface status.

```

user@host# run show interfaces pp0
Physical interface: pp0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 510
  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 packets : 0
Output packets: 0

Logical interface pp0.0 (Index 70) (SNMP ifIndex 522)
Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: PPPoE
PPPoE:
  State: SessionUp, Session ID: 8,
  Session AC name: belur, Remote MAC address: 00:90:1a:43:18:d1,
  Configured AC name: None, Service name: None,
  Auto-reconnect timeout: 120 seconds, Idle timeout: Never,
  Underlying interface: pt-1/0/0.0 (Index 69)
Input packets : 12
Output packets: 11
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive: Input: 0 (never), Output: 4 (00:00:03 ago)
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mp1s:
Not-configured
CHAP state: Success
PAP state: Closed
Security: Zone: Null
Protocol inet, MTU: 1474
  Flags: Negotiate-Address
  Addresses, Flags: Kernel Is-Preferred Is-Primary
  Destination: 12.12.12.1, Local: 12.12.12.12

```

2. Verify the end-to-end data path connectivity.

```

user@host# run show interfaces pt-1/0/0 terse
Interface      Admin Link Proto  Local          Remote
pt-1/0/0       up    up
pt-1/0/0.0     up    up

[edit]
user@host# run show interfaces pp0 terse
Interface      Admin Link Proto  Local          Remote
pp0            up    up
pp0.0          up    up  inet  12.12.12.12    --> 12.12.12.1

[edit]
user@host# run ping 12.12.12.1 count 100 rapid
PING 12.12.12.1 (12.12.12.1): 56 data bytes
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
--- 12.12.12.1 ping statistics ---
100 packets transmitted, 100 packets received, 0% packet loss
round-trip min/avg/max/stddev = 16.168/17.452/23.299/2.016 ms

```

- Related Documentation**
- [VDSL2 Interface Technology Overview on page 75](#)
 - [Example: Configuring VDSL2 Interfaces \(Basic\) on page 81](#)
 - [Example: Configuring VDSL2 Interfaces in ADSL Mode \(Detail\)](#)

Upgrading the VDSL PIC Firmware

Supported Platforms [SRX320](#), [SRX340](#), [SRX345](#)

This topic shows how to upgrade the VDSL PIC firmware on SRX Series devices.

Before you begin:

Check the current firmware version of the VDSL PIC.

```
user@host> show system firmware
Part Type Tag Current Available Status version
FPC 1
PIC 0 VDSLBCM 10 2.10.0 OK
Routing Engine 0 RE BIOS 0 2.0 OK
Routing Engine 0 RE BIOS Backup 1 2.0 OK
Routing Engine 0 RE FPGA 14 1.0.0 OK
```

This section describes the step-by-step procedure to upgrade VDSL PIC firmware.

1. Mount or copy the firmware package to the SRX Series device.

If the file has been obtained from JTAC, use FTP or SCP to load the firmware file on the device. Save the file in the `/var/tmp` directory.

2. Upgrade the firmware on the SRX Series device.

To install the firmware package on the device and make it available for upgrading, use the following command:

```
user@host> request system software add no-copy no-validate
jfirmware-srxsme-11.4R2.7-signed.tgz
```

3. To check if the firmware package is available on the SRX Series device, use the following command:

```
user@host> show version

Hostname: user

Model: srx210h

JUNOS Software Release [12.1I20120123_0941]

JUNOS Firmware Software Suite [11.4R2.7]
```

4. To verify the VDSL PIM slot, use the following command:

```
user@host> show chassis hardware
```

5. To initiate a firmware upgrade, use the following command:

```
user@host> request system firmware upgrade pic fpc-slot <no.> pic-slot 0 tag 10
```

6. To check the status of the upgraded firmware, use the following command:

```
user@host> show system firmware
Part Type Tag Current Available Status version
FPC 1
PIC 0 VDSLBCM 10 2.10.0 2.11.0
Routing Engine 0 RE BIOS 0 2.0 OK
Routing Engine 0 RE BIOS Backup 1 2.0 OK
Routing Engine 0 RE FPGA 14 203.0.113.45.0.0 OK
```

7. To enable the upgraded firmware, restart the FPC slot in which the VDSL PIM is installed.

```
user@host> restart fpc <no.>
```

FPC 1 restarted

8. To verify the firmware upgrade is complete, use the following command:

```
user@host> show system firmware
```

```
Part Type Tag Current Available Status version
FPC 1
PIC 0 VDSLBCM 10 2.11.0 2.11.0 OK
Routing Engine 0 RE BIOS 0 2.0 OK
Routing Engine 0 RE BIOS Backup 1 2.0 OK
Routing Engine 0 RE FPGA 14 203.0.113.45.0.0 OK
```

PART 4

Configuring Ethernet Interfaces

- [Performing Initial Configuration on Ethernet Interfaces on page 117](#)
- [Configuring Aggregated Ethernet Interfaces on page 133](#)
- [Configuring Link Aggregation Control Protocol on page 149](#)
- [Configuring Gigabit Ethernet Physical Interface Modules on page 165](#)
- [Configuring Ethernet OAM Link Fault Management on page 181](#)
- [Configuring Power over Ethernet on page 187](#)

CHAPTER 8

Performing Initial Configuration on Ethernet Interfaces

- [Understanding Ethernet Interfaces on page 117](#)
- [Understanding Static ARP Entries on Ethernet Interfaces on page 121](#)
- [Understanding Promiscuous Mode on Ethernet Interface on page 121](#)
- [Example: Creating an Ethernet Interface on page 122](#)
- [Example: Deleting an Ethernet Interface on page 123](#)
- [Example: Configuring Static ARP Entries on Ethernet Interfaces on page 124](#)
- [Enabling and Disabling Promiscuous Mode on Ethernet Interfaces \(CLI Procedure\) on page 127](#)
- [Example: Configuring Promiscuous Mode on the SRX5K-MPC on page 127](#)

Understanding Ethernet Interfaces

Supported Platforms [SRX Series, vSRX](#)

Ethernet is a Layer 2 technology that operates in a shared bus topology. Ethernet supports broadcast transmission, uses best-effort delivery, and has distributed access control. Ethernet is a point-to-multipoint technology.

In a shared bus topology, all devices connect to a single, shared physical link through which all data transmissions are sent. All traffic is broadcast so that all devices within the topology receive every transmission. The devices within a single Ethernet topology make up a broadcast domain.

Ethernet uses best-effort delivery to broadcast traffic. The physical hardware provides no information to the sender about whether the traffic was received. If the receiving host is offline, traffic to the host is lost. Although the Ethernet data link protocol does not inform the sender about lost packets, higher layer protocols such as TCP/IP might provide this type of notification.

This topic contains the following sections:

- [Ethernet Access Control and Transmission on page 118](#)
- [Collisions and Detection on page 118](#)

- [Collision Domains and LAN Segments on page 119](#)
- [Broadcast Domains on page 120](#)
- [Ethernet Frames on page 120](#)

Ethernet Access Control and Transmission

Ethernet's access control is distributed because Ethernet has no central mechanism that grants access to the physical medium within the network. Instead, Ethernet uses carrier-sense multiple access with collision detection (CSMA/CD). Because multiple devices on an Ethernet network can access the physical medium, or wire, simultaneously, each device must determine whether the physical medium is in use. Each host listens on the wire to determine if a message is being transmitted. If it detects no transmission, the host begins transmitting its own data.

The length of each transmission is determined by fixed Ethernet packet sizes. By fixing the length of each transmission and enforcing a minimum idle time between transmissions, Ethernet ensures that no pair of communicating devices on the network can monopolize the wire and block others from sending and receiving traffic.

Collisions and Detection

When a device on an Ethernet network begins transmitting data, the data takes a finite amount of time to reach all hosts on the network. Because of this delay, or latency, in transmitting traffic, a device might detect an idle state on the wire just as another device initially begins its transmission. As a result, two devices might send traffic across a single wire at the same time. When the two electrical signals collide, they become scrambled so that both transmissions are effectively lost.

Collision Detection

To handle collisions, Ethernet devices monitor the link while they are transmitting data. The monitoring process is known as collision detection. If a device detects a foreign signal while it is transmitting, it terminates the transmission and attempts to transmit again only after detecting an idle state on the wire. Collisions continue to occur if two colliding devices both wait the same amount of time before retransmitting. To avoid this condition, Ethernet devices use a binary exponential backoff algorithm.

Backoff Algorithm

With the binary exponential backoff algorithm, each device that sends a colliding transmission randomly selects a value within a range. The value represents the number of transmission times that the device must wait before retransmitting its data. If another collision occurs, the range of values is doubled and retransmission takes place again. Each time a collision occurs, the range of values doubles, to reduce the likelihood that two hosts on the same network can select the same retransmission time.

[Table 17 on page 118](#) shows collision rounds up to round 10.

Table 17: Collision Backoff Algorithm Rounds

Round	Size of Set	Elements in the Set
1	2	{0,1}

Table 17: Collision Backoff Algorithm Rounds (*continued*)

Round	Size of Set	Elements in the Set
2	4	{0,1,2,3}
3	8	{0,1,2,3,....,7}
4	16	{0,1,2,3,4,....,15}
5	32	{0,1,2,3,4,5,....,31}
6	64	{0,1,2,3,4,5,6,....,63}
7	128	{0,1,2,3,4,5,6,7,....,127}
8	256	{0,1,2,3,4,5,6,7,8,....,255}
9	512	{0,1,2,3,4,5,6,7,8,9,....,511}
10	1024	{0,1,2,3,4,5,6,7,8,9,10,....,1023}

Collision Domains and LAN Segments

Collisions are confined to a physical wire over which data is broadcast. Because the physical wires are subject to signal collisions, individual LAN segments are known as *collision domains*. Although the physical limitations on the length of an Ethernet cable restrict the length of a LAN segment, multiple collision domains can be interconnected by repeaters, bridges, and switches.

Repeaters

Repeaters are electronic devices that act on analog signals. Repeaters relay all electronic signals from one wire to another. A single repeater can double the distance between two devices on an Ethernet network. However, the Ethernet specification restricts the number of repeaters between any two devices on an Ethernet network to two, because collision detection with latencies increases in complexity as the wire length and number of repeaters increase.

Bridges and Switches

Bridges and switches combine LAN segments into a single Ethernet network by using multiple ports to connect the physical wires in each segment. Although bridges and switches are fundamentally the same, bridges generally provide more management and more interface ports. As Ethernet packets flow through a bridge, the bridge tracks the source MAC address of the packets and stores the addresses and their associated input ports in an interface table. As it receives subsequent packets, the bridge examines its interface table and takes one of the following actions:

- If the destination address does not match an address in the interface table, the bridge transmits the packet to all hosts on the network using the Ethernet broadcast address.

- If the destination address maps to the port through which the packet was received, the bridge or switch discards the packet. Because the other devices on the LAN segment also received the packet, the bridge does not need to retransmit it.
- If the destination address maps to a port other than the one through which the packet was received, the bridge transmits the packet through the appropriate port to the corresponding LAN segment.

Broadcast Domains

The combination of all the LAN segments within an Ethernet network is called a *broadcast domain*. In the absence of any signaling devices such as a repeater, bridge, or switch, the broadcast domain is simply the physical wire that makes up the connections in the network. If a bridge or switch is used, the broadcast domain consists of the entire LAN.

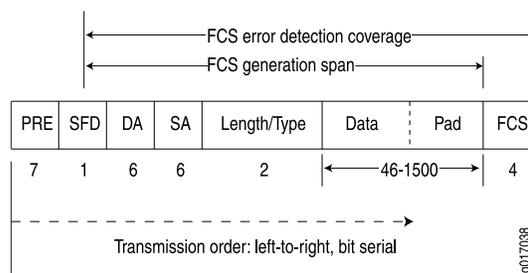


NOTE: On all branch SRX Series devices, the subnet directed broadcast feature is not supported.

Ethernet Frames

Data is transmitted through an Ethernet network in frames. The frames are of variable length, ranging from 64 octets to 1518 octets, including the header, payload, and cyclic redundancy check (CRC) value. [Figure 10 on page 120](#) shows the Ethernet frame format.

Figure 10: Ethernet Frame Format



Ethernet frames have the following fields:

- The preamble (PRE) field is 7 octets of alternating 0s and 1s. The predictable format in the preamble allows receiving interfaces to synchronize themselves to the data being sent. The preamble is followed by a 1-octet start-of-frame delimiter (SFD).
- The destination address (DA) and source address (SA) fields contain the 6-octet (48-bit) MAC addresses for the destination and source ports on the network. These Layer 2 addresses uniquely identify the devices on the LAN.
- The Length/Type field is a 2-octet field that either indicates the length of the frame's data field or identifies the protocol stack associated with the frame. Here are some common frame types:
 - AppleTalk—**0x809B**
 - AppleTalk ARP—**0x80F3**

- DECnet—0x6003
- IP—0x0800
- IPX—0x8137
- Loopback—0x9000
- XNS—0x0600
- The Data field contains the packet payload.
- The frame check sequence (FCS) is a 4-octet field that contains the calculated CRC value. This value is calculated by the originating host and appended to the frame. When it receives the frames, the receiving host calculates the CRC and checks it against this appended value to verify the integrity of the received frame.

Related Documentation

- [Understanding Interfaces on page 3](#)
- [Example: Creating an Ethernet Interface on page 122](#)
- [Example: Deleting an Ethernet Interface on page 123](#)
- [Understanding Static ARP Entries on Ethernet Interfaces on page 121](#)
- [Understanding Promiscuous Mode on Ethernet Interface on page 121](#)

Understanding Static ARP Entries on Ethernet Interfaces

Supported Platforms SRX Series, vSRX

By default, the device responds to an Address Resolution Protocol (ARP) request only if the destination address of the ARP request is on the local network of the incoming interface. For Fast Ethernet or Gigabit Ethernet interfaces, you can configure static ARP entries that associate the IP addresses of nodes on the same Ethernet subnet with their media access control (MAC) addresses. These static ARP entries enable the device to respond to ARP requests even if the destination address of the ARP request is not local to the incoming Ethernet interface.

Related Documentation

- [Understanding Ethernet Interfaces on page 117](#)
- [Example: Configuring Static ARP Entries on Ethernet Interfaces on page 124](#)

Understanding Promiscuous Mode on Ethernet Interface

Supported Platforms SRX1500, SRX5400, SRX5600, SRX5800, vSRX

When promiscuous mode is enabled on a Layer 3 Ethernet interface, all packets received on the interface are sent to the central point or Services Processing Unit (SPU) regardless of the destination MAC address of the packet. You can also enable promiscuous mode on chassis cluster redundant Ethernet interfaces and aggregated Ethernet interfaces. If you enable promiscuous mode on a redundant Ethernet interface, promiscuous mode is then enabled on any child physical interfaces. If you enable promiscuous mode on an

aggregated Ethernet interface, promiscuous mode is then enabled on all member interfaces.

Understanding Promiscuous Mode on the SRX5K-MPC

The promiscuous mode function is supported on 1-Gigabit, 10-Gigabit, 40-Gigabit, and 100-Gigabit Ethernet interfaces on the I/O cards (IOCs) and the SRX5000 line Module Port Concentrator (SRX5K-MPC).

When promiscuous mode is enabled on a Layer 3 Ethernet interface, all packets received on the interface are sent to the central point or to the Services Processing Unit (SPU) regardless of the destination MAC address of the packet.

By default, an interface enables MAC filtering. You can configure promiscuous mode on the interface to disable MAC filtering. When you delete the promiscuous mode configuration, the interface will perform MAC filtering again.

You can change the MAC address of an interface even when the interface is operating in promiscuous mode. When the interface is operating in normal mode again, the MAC filtering function on the IOC uses the new MAC address to filter the packets.

You can also enable promiscuous mode on chassis cluster redundant Ethernet interfaces and aggregated Ethernet interfaces. If you enable promiscuous mode on a redundant Ethernet interface, promiscuous mode is then enabled on any child physical interfaces. If you enable promiscuous mode on an aggregated Ethernet interface, promiscuous mode is then enabled on all member interfaces.

Related Documentation

- [Understanding Ethernet Interfaces on page 117](#)
- [Enabling and Disabling Promiscuous Mode on Ethernet Interfaces \(CLI Procedure\) on page 127](#)
- [Example: Configuring Promiscuous Mode on the SRX5K-MPC on page 127](#)

Example: Creating an Ethernet Interface

Supported Platforms [SRX Series, vSRX](#)

This example shows how to create an Ethernet interface.

- [Requirements on page 122](#)
- [Overview on page 123](#)
- [Configuration on page 123](#)

Requirements

No special configuration beyond device initialization is required before configuring an interface.

Overview

In this example, you create the ge-1/0/0 Ethernet interface and set the logical interface to 0. The logical unit number can range from 0 to 16,384. You can also add values for properties that you need to configure on the logical interface, such as logical encapsulation or protocol family.

Configuration

Step-by-Step Procedure

To configure an Ethernet interface:

1. Create the Ethernet interface and set the logical interface.


```
[edit]
user@host# edit interfaces ge-1/0/0 unit 0
```
2. If you are done configuring the device, commit the configuration.


```
[edit]
user@host# commit
```

Verification

Purpose Verify if the configuration is working properly after creating the interface.

Action From operational mode, enter the **show interfaces** command.

Related Documentation

- [Understanding Ethernet Interfaces on page 117](#)
- [Example: Deleting an Ethernet Interface on page 123](#)

Example: Deleting an Ethernet Interface

Supported Platforms SRX Series, vSRX

This example shows how to delete an Ethernet interface.

- [Requirements on page 123](#)
- [Overview on page 123](#)
- [Configuration on page 124](#)

Requirements

No special configuration beyond device initialization is required before configuring an interface.

Overview

In this example, you delete the ge-1/0/0 interface.



NOTE: Performing this action removes the interface from the software configuration and disables it. Network interfaces remain physically present, and their identifiers continue to appear on J-Web pages.

Configuration

Step-by-Step Procedure

To delete an Ethernet interface:

1. Specify the interface you want to delete.


```
[edit]
user@host# delete interfaces ge-1/0/0
```
2. If you are done configuring the device, commit the configuration.


```
[edit]
user@host# commit
```

Verification

Purpose Verify if the configuration is working properly after deleting the interface.

Action From operational mode, enter the **show interfaces** command.

Related Documentation

- [Understanding Ethernet Interfaces on page 117](#)
- [Example: Creating an Ethernet Interface on page 122](#)

Example: Configuring Static ARP Entries on Ethernet Interfaces

Supported Platforms SRX Series, vSRX

- [Requirements on page 124](#)
- [Overview on page 124](#)
- [Configuration on page 125](#)
- [Verification on page 125](#)

Requirements

No special configuration beyond device initialization is required before creating an interface.

Overview

In this example, you configure a static ARP entry on the logical unit 0 of the ge-0/0/3 Gigabit Ethernet interface. The entry consists of the interface's IP address (10.1.1.1/24) and the corresponding MAC address of a node on the same Ethernet subnet (00:ff:85:7f:78:03). The example also configures the device to reply to ARP requests from the node using the publish option.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/3 unit 0 family inet address 10.1.1.1/24 arp 10.1.1.3 mac
00:ff:85:7f:78:03
set interfaces ge-0/0/3 unit 0 family inet address 10.1.1.1/24 arp 10.1.1.3 publish
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a static ARP entry on an Ethernet interface:

1. Create the Gigabit Ethernet interface.

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

2. Configure a static ARP entry.

```
[edit interfaces ge-0/0/3]
user@host# edit unit 0 family inet address 10.1.1.1/24
```

3. Set the IP address of the subnet node and the corresponding MAC address.

```
[edit interfaces ge-0/0/3 unit 0 family inet address 10.1.1.1/24]
user@host# set arp 10.1.1.3 mac 00:ff:85:7f:78:03 publish
```

Results From configuration mode, confirm your configuration by entering the **show interfaces ge-0/0/3** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host#show interfaces ge-0/0/3
unit 0 {
  family inet {
    address 10.1.1.1/24 {
      arp 10.1.1.3 mac 00:ff:85:7f:78:03 publish;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying Static ARP Configurations on page 126](#)
- [Verifying the Link State of All Interfaces on page 126](#)
- [Verifying Interface Properties on page 126](#)

Verifying Static ARP Configurations

Purpose Verify the IP address and MAC (hardware) address of the node.

Action From operational mode, enter the **show interfaces ge-0/0/3** command.

Verifying the Link State of All Interfaces

Purpose Verify that all interfaces on the device are operational using the ping tool on each peer address in the network.

Action For each interface on the device:

1. In the J-Web interface, select **Troubleshoot>Ping Host**.
2. In the Remote Host box, type the address of the interface for which you want to verify the link state.
3. Click **Start**. The output appears on a separate page.

```
PING 10.10.10.10 : 56 data bytes
64 bytes from 10.10.10.10: icmp_seq=0 ttl=255 time=0.382 ms
64 bytes from 10.10.10.10: icmp_seq=1 ttl=255 time=0.266 ms
```

If the interface is operational, it generates an ICMP response. If this response is received, the round-trip time in milliseconds is listed in the time field..

Verifying Interface Properties

Purpose Verify that the interface properties are correct.

Action From operational mode, enter the **show interfaces detail** command.

```
user@host> show interfaces detail
Physical interface: ge-0/0/3, Enabled, Physical link is Up
  Interface index: 134, SNMP ifIndex: 27, Generation: 17
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps 16384
  Link flags     : None
  CoS queues    : 4 supported
  Hold-times    : Up 0 ms, Down 0 ms
  Current address: 00:90:69:87:44:9d, Hardware address: 00:90:69:87:44:9d
  Last flapped  : 2004-08-25 15:42:30 PDT (4w5d 22:49 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
  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

```

The output shows a summary of interface information. Verify the following information:

- The physical interface is Enabled. If the interface is shown as Disabled, do one of the following:
 - In the CLI configuration editor, delete the **disable** statement at the [edit interfaces ge-0/0/3] level of the configuration hierarchy.
 - In the J-Web configuration editor, clear the **Disable** check box on the Interfaces > ge-0/0/3 page.
- The physical link is Up. A link state of Down indicates a problem with the interface module, interface port, or physical connection (link-layer errors).
- The Last Flapped time is an expected value. The Last Flapped time indicates the last time the physical interface became unavailable and then available again. Unexpected flapping indicates likely link-layer errors.
- The traffic statistics reflect expected input and output rates. Verify that the number of inbound and outbound bytes and packets matches expected throughput for the physical interface. To clear the statistics and see only new changes, use the **clear interfaces statistics ge-0/0/3** command.

Related Documentation

- [Understanding Static ARP Entries on Ethernet Interfaces on page 121](#)

Enabling and Disabling Promiscuous Mode on Ethernet Interfaces (CLI Procedure)

Supported Platforms [SRX1400](#), [SRX3400](#), [SRX3600](#), [SRX5600](#), [SRX5800](#)

To enable promiscuous mode on an interface:

```
user@host# set interfaces interface-name promiscuous-mode
```

To disable promiscuous mode on an interface:

```
user@host# delete interfaces interface-name promiscuous-mode
```

Related Documentation

- [Understanding Promiscuous Mode on Ethernet Interface on page 121](#)
- [Understanding Ethernet Interfaces on page 117](#)

Example: Configuring Promiscuous Mode on the SRX5K-MPC

Supported Platforms [SRX5400](#), [SRX5600](#), [SRX5800](#)

This example shows how to configure promiscuous mode on an SRX5K-MPC interface in an SRX5600 to disable MAC address filtering.

- [Requirements on page 128](#)
- [Overview on page 128](#)
- [Configuration on page 128](#)
- [Verification on page 129](#)

Requirements

This example uses the following hardware and software components:

- An SRX5600 with an SRX5K-MPC that includes a 100-Gigabit Ethernet CFP transceiver
- Junos OS Release 12.1X47-D10 or later

No special configuration beyond device initialization is required before configuring this feature.

Overview

By default, the interfaces on an SRX5K-MPC have MAC address filtering enabled. In this example, you configure promiscuous mode on an interface to disable MAC address filtering. Then you delete promiscuous mode to reenables MAC address filtering on the interface.

Configuration

Configuring Promiscuous Mode on an Interface

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces et-4/0/0 unit 0 family inet address 10.1.1.1/24
set interfaces et-4/0/0 promiscuous-mode
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see the [Junos OS CLI User Guide](#).

To configure promiscuous mode:

1. Configure the ingress interface.

```
[edit interfaces]
user@host# set et-4/0/0 unit 0 family inet address 10.1.1.1/24
```

2. Enable promiscuous mode on the interface.

```
[edit interfaces]
user@host# set et-4/0/0 promiscuous-mode
```

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

```
[edit]
user@host# show interfaces
et-4/0/0 {
  promiscuous-mode;
  unit 0 {
    family inet {
      address 10.1.1.1/24;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Disabling Promiscuous Mode on an Interface

CLI Quick Configuration To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
user@host# delete interfaces et-4/0/0 promiscuous-mode
```

Step-by-Step Procedure To disable promiscuous mode:

1. Disable promiscuous mode on the interface.

```
[edit]
user@host# delete interfaces et-4/0/0 promiscuous-mode
```

Verification

Confirm that the configuration is working properly.

- [Verifying That Promiscuous Mode Is Enabled on the SRX5K-MPC on page 129](#)
- [Verifying the Status of Promiscuous Mode on page 130](#)
- [Verifying That Promiscuous Mode Is Disabled on page 131](#)

Verifying That Promiscuous Mode Is Enabled on the SRX5K-MPC

Purpose Verify that promiscuous mode is enabled on the interface.

Action From operational mode, enter the **monitor interface traffic** command.

```
user@host> monitor interface traffic

Physical interface: et-4/0/0, Enabled, Physical link is Up
Interface index: 137, SNMP ifIndex: 511
Link-level type: Ethernet, MTU: 1518, Speed: 100Gbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled
Device flags   : Present Running
```

```

Interface flags: Promiscuous SNMP-Traps Internal: 0x4000
CoS queues      : 8 supported, 8 maximum usable queues
Current address: 2c:21:72:3a:05:28, Hardware address: 2c:21:72:3a:05:28
Last flapped   : 2014-01-17 14:44:53 PST (5d 06:30 ago)
Input rate     : 0 bps (0 pps)
Output rate    : 0 bps (0 pps)
Active alarms  : None
Active defects : None
PCS statistics                               Seconds
  Bit errors                                0
  Errored blocks                            0
    
```

```

Logical interface et-4/0/0.0 (Index 71) (SNMP ifIndex 513)
Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1351 ] Encapsulation: ENET2
Input packets : 0
Output packets: 0
Security: Zone: HOST
Allowed host-inbound traffic : any-service bfd bgp dvmrp igmp ldp msdp nhrp
ospf pgm pim rip router-discovery rsvp sap vrrp
Protocol inet, MTU: 1500
  Flags: Sendbroadcast-pkt-to-re
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 122.122.122/24, Local: 122.122.122.1,
    Broadcast: 122.122.122.255
Protocol multiservice, MTU: Unlimited
  Flags: Is-Primary
    
```

```

Logical interface et-4/0/0.32767 (Index 72) (SNMP ifIndex 517)
Flags: SNMP-Traps 0x40040000 VLAN-Tag [ 0x0000.0 ] Encapsulation: ENET2
Input packets : 0
Output packets: 0
Security: Zone: HOST
Allowed host-inbound traffic : any-service bfd bgp dvmrp igmp ldp msdp nhrp
ospf pgm pim rip router-discovery rsvp sap vrrp
Protocol multiservice, MTU: Unlimited
  Flags: None
    
```

Meaning The **Interface flags: Promiscuous** field shows that promiscuous mode is enabled on the interface.

Verifying the Status of Promiscuous Mode

Purpose Verify that promiscuous mode works on the **et-4/0/0** interface.

Action Send traffic into the **et-4/0/0** interface with a MAC address that is different from the interface MAC address and turn on promiscuous mode.

From operational mode, enter the **monitor interface traffic** command.

```
user@host> monitor interface traffic
```

Interface	Link	Input packets	(pps)	Output packets	(pps)
gr-0/0/0	Up	0	(0)	0	(0)
ip-0/0/0	Up	0	(0)	0	(0)
lt-0/0/0	Up	0	(0)	0	(0)
xe-1/2/0	Down	0	(0)	0	(0)
xe-1/2/1	Down	0	(0)	0	(0)
xe-1/2/2	Down	0	(0)	0	(0)
xe-1/2/3	Down	0	(0)	0	(0)

xe-1/2/4	Down	0	(0)	0	(0)
xe-1/2/5	Down	0	(0)	0	(0)
xe-1/2/6	Down	0	(0)	0	(0)
xe-1/2/7	Down	0	(0)	0	(0)
xe-1/2/8	Down	0	(0)	0	(0)
xe-1/2/9	Down	0	(0)	0	(0)
et-4/0/0	Up	4403996	(100002)	0	(0)
et-4/2/0	Up	3	(0)	4403924	(99997)
avs0	Up	0	(0)	0	(0)
avs1	Up	0	(0)	0	(0)
dsc	Up	0		0	
em0	Up	15965		14056	

Meaning The **input packets** and **pps** fields show that traffic is passing through the **et-4/0/0** interface as expected after promiscuous mode is enabled.

Verifying That Promiscuous Mode Is Disabled

Purpose Verify that disabled promiscuous mode works on the **et-4/0/0** interface.

Action Send traffic into the **et-4/0/0** interface with a MAC address that is different from the interface MAC address and turn off promiscuous mode.

From operational mode, enter the **monitor interface traffic** command.

```
user@host> monitor interface traffic
```

Interface	Link	Input packets	(pps)	Output packets	(pps)
gr-0/0/0	Up	0	(0)	0	(0)
ip-0/0/0	Up	0	(0)	0	(0)
lt-0/0/0	Up	0	(0)	0	(0)
xe-1/2/0	Down	0	(0)	0	(0)
xe-1/2/1	Down	0	(0)	0	(0)
xe-1/2/2	Down	0	(0)	0	(0)
xe-1/2/3	Down	0	(0)	0	(0)
xe-1/2/4	Down	0	(0)	0	(0)
xe-1/2/5	Down	0	(0)	0	(0)
xe-1/2/6	Down	0	(0)	0	(0)
xe-1/2/7	Down	0	(0)	0	(0)
xe-1/2/8	Down	0	(0)	0	(0)
xe-1/2/9	Down	0	(0)	0	(0)
et-4/0/0	Up	11505495	(0)	0	(0)
et-4/2/0	Up	6	(0)	11505425	(0)
avs0	Up	0	(0)	0	(0)
avs1	Up	0	(0)	0	(0)
dsc	Up	0		0	
em0	Up	37964		31739	

Meaning The **pps** field shows that the traffic is not passing through the **et-4/0/0** interface after promiscuous mode is disabled.

Related Documentation

- [Understanding Promiscuous Mode on Ethernet Interface on page 121](#)
- [Enabling and Disabling Promiscuous Mode on Ethernet Interfaces \(CLI Procedure\) on page 127](#)

CHAPTER 9

Configuring Aggregated Ethernet Interfaces

- [Understanding Aggregated Ethernet Interfaces on page 133](#)
- [Aggregated Ethernet Interfaces Configuration Overview on page 136](#)
- [Understanding the Aggregated Ethernet Interfaces Device Count on page 136](#)
- [Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device on page 137](#)
- [Understanding Physical Interfaces for Aggregated Ethernet Interfaces on page 138](#)
- [Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces on page 138](#)
- [Understanding Aggregated Ethernet Interface Link Speed on page 139](#)
- [Example: Configuring Aggregated Ethernet Link Speed on page 140](#)
- [Understanding Minimum Links for Aggregated Ethernet Interfaces on page 141](#)
- [Example: Configuring Aggregated Ethernet Minimum Links on page 141](#)
- [Understanding Aggregated Ethernet Interface Removal on page 142](#)
- [Example: Deleting Aggregated Ethernet Interfaces on page 143](#)
- [Example: Deleting Aggregated Ethernet Interface Contents on page 144](#)
- [Verifying Aggregated Ethernet Interfaces on page 145](#)
- [Understanding VLAN Tagging for Aggregated Ethernet Interfaces on page 146](#)
- [Understanding Promiscuous Mode for Aggregated Ethernet Interfaces on page 147](#)

Understanding Aggregated Ethernet Interfaces

Supported Platforms SRX5400, SRX5600, SRX5800

Link aggregation of Ethernet interfaces is defined in the IEEE 802.3ad standard. Junos OS implementation of 802.3ad balances traffic across the member links within an aggregated Ethernet bundle based on Layer 3 information carried in the packet, Layer 4 information carried in the packet, or both, or based on session ID data. (The session ID data has higher precedence than the Layer 3 or 4 information.) This implementation uses the same load-balancing algorithm used for per-packet load balancing.

Aggregated Ethernet interfaces can be Layer 3 interfaces (VLAN-tagged or untagged) and Layer 2 interfaces.



NOTE: This topic is specific to the SRX3000 and SRX5000 line devices. For information about link aggregation for other SRX Series devices, see the *Layer 2 Bridging and Transparent Mode for Security Devices*.

This topic contains the following sections:

- [LAGs on page 134](#)
- [LACP on page 134](#)

LAGs

You can combine multiple physical Ethernet ports to form a logical point-to-point link, known as a link aggregation group (LAG) or bundle, such that a media access control (MAC) client can treat the LAG as if it were a single link. Support for LAGs based on IEEE 802.3ad makes it possible to aggregate physical interface links on your device. LAGs provide increased interface bandwidth and link availability by linking physical ports and load-balancing traffic crossing the combined interface. For the LAG to operate correctly, it is necessary to coordinate the two end systems connected by the LAG, either manually or automatically.

Internally, a LAG is a virtual interface presented on SRX3000 and SRX5000 line devices or on any system (consisting of devices such as routers and switches) supporting 802.3ad link aggregation. Externally, a LAG corresponds to a bundle of physical Ethernet links connected between an SRX3000 or SRX5000 line device and another system capable of link aggregation. This bundle of physical links is a virtual link.

Follow these guidelines for aggregated Ethernet support for the SRX3000 and SRX5000 lines:

- The devices support a maximum of 16 physical interfaces per single aggregated Ethernet bundle.
- Aggregated Ethernet interfaces can use interfaces from the same or different Flexible PIC Concentrators (FPCs) and PICs.
- On the aggregated bundle, capabilities such as MAC accounting, VLAN rewrites, and VLAN queuing are available.

LACP

Junos OS supports the Link Aggregation Control Protocol (LACP), which is a subcomponent of IEEE 802.3ad. LACP provides additional functionality for LAGs.

Starting with Junos OS Release 15.1X49-D40, LACP is supported on Layer 2 transparent mode in addition to existing support on Layer 3 mode. For information about link aggregation for other SRX Series devices, see the *Layer 2 Bridging and Transparent Mode for Security Devices*.

LACP provides a standardized means for exchanging information between partner (remote or far-end of the link) systems on a link. This exchange allows their link aggregation control instances to reach agreement on the identity of the LAG to which the link belongs, and then to move the link to that LAG. This exchange also enables the transmission and reception processes for the link to function in an orderly manner.

For example, when LACP is not enabled, a local LAG might attempt to transmit packets to a remote individual interface, which causes the communication to fail. (An individual interface is a nonaggregatable interface.) When LACP is enabled, a local LAG cannot transmit packets unless a LAG with LACP is also configured on the remote end of the link.

You configure an aggregated Ethernet virtual link by specifying the link number as a physical device. Then you associate a set of ports that have the same speed and are in full-duplex mode. The physical ports can be 100-megabit Ethernet, 1-Gigabit Ethernet, and 10-Gigabit Ethernet.

When configuring LACP, follow these guidelines:

- LACP does not support automatic configuration on SRX3000 and SRX5000 line devices, but partner systems are allowed to perform automatic configuration. When an SRX3000 or SRX5000 line device is connected to a fully 802.3ad-compliant partner system, static configuration of LAGs is initiated on the SRX3000 and SRX5000 line device side, and static configuration is not needed on the partner side.
- When an SRX3000 or SRX5000 line device is connected to a Juniper Networks MX Series router, static configuration of LAGs is needed at both the actor (local or near-end of the link) and partner systems.
- Although the LACP functions on the SRX3000 and SRX5000 line devices are similar to the LACP features on Juniper Networks MX Series routers, the following LACP features on MX Series routers are not supported on SRX3000 and SRX5000 line devices: link protection, system priority, and port priority for aggregated Ethernet interfaces. Instead, SRX3000 and SRX5000 line devices provide active/standby support with redundant Ethernet interface LAGs in chassis cluster deployments.

LACP is supported in standalone deployments, where aggregated Ethernet interfaces are supported, and in chassis cluster deployments, where aggregated Ethernet interfaces and redundant Ethernet interfaces are supported simultaneously.

**Related
Documentation**

- [Understanding Ethernet Interfaces on page 117](#)
- [Aggregated Ethernet Interfaces Configuration Overview on page 136](#)
- [Understanding LACP on Standalone Devices on page 149](#)
- [Understanding LACP on Chassis Clusters on page 153](#)
- [Understanding VLAN Tagging for Aggregated Ethernet Interfaces on page 146](#)
- [Understanding Promiscuous Mode for Aggregated Ethernet Interfaces on page 147](#)

Aggregated Ethernet Interfaces Configuration Overview

Supported Platforms [SRX Series](#)



NOTE: This topic is specific to the SRX3000 and SRX5000 line devices.

To configure an aggregated Ethernet interface:

1. Set the number of aggregated Ethernet interfaces on the device. See [“Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device”](#) on page 137.
2. Associate a physical interface with the aggregated Ethernet interface. See [“Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces”](#) on page 138.
3. (Optional) Set the required link speed for all the interfaces included in the bundle. See [“Example: Configuring Aggregated Ethernet Link Speed”](#) on page 140.
4. (Optional) Configure the minimum number of links that must be up for the bundle as a whole to be labeled as up. See [“Example: Configuring Aggregated Ethernet Minimum Links”](#) on page 141.
5. (Optional) Enable or disable VLAN tagging. See [“Understanding VLAN Tagging for Aggregated Ethernet Interfaces”](#) on page 146.
6. (Optional) Enable promiscuous mode. See [“Understanding Promiscuous Mode for Aggregated Ethernet Interfaces”](#) on page 147.

Related Documentation

- [Layer 2 Bridging and Transparent Mode for Security Devices](#)
- [Understanding Aggregated Ethernet Interfaces on page 133](#)
- [Example: Configuring LACP on Standalone Devices on page 150](#)
- [Example: Configuring LACP on Chassis Clusters on page 155](#)

Understanding the Aggregated Ethernet Interfaces Device Count

Supported Platforms [SRX Series](#)

By default, no aggregated Ethernet interfaces are created. You must set the number of aggregated Ethernet interfaces on the routing device before you can configure them. Once you set the device count, the system creates that number of empty aggregated Ethernet interfaces. A globally unique MAC address is assigned to every aggregated Ethernet interface. More aggregated Ethernet interfaces can be created by increasing the parameter.

The maximum number of aggregated devices you can configure is 128. The aggregated interfaces are numbered from ae0 through ae127.

Similarly, you can permanently remove an aggregated Ethernet interface from the device configuration by deleting it from the device count. When you reduce the device count,

only the aggregated Ethernet interface objects at the end of the list are removed, leaving the newly specified number of interfaces. That is, if you set the device count to 10 and then reduce it to 6, the system removes the last 4 interface objects from the list.



WARNING: Be aware that this approach deletes the aggregated Ethernet interface and *all* of its objects from the device configuration.

Related Documentation

- [Understanding Aggregated Ethernet Interfaces on page 133](#)
- [Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device on page 137](#)
- [Example: Deleting Aggregated Ethernet Interfaces on page 143](#)

Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device

Supported Platforms SRX Series

This example shows how to configure the number of aggregated Ethernet interfaces on a device.

- [Requirements on page 137](#)
- [Overview on page 137](#)
- [Configuration on page 137](#)
- [Verification on page 138](#)

Requirements

No special configuration beyond device initialization is required before configuring an interface.

Overview

In this example, you create two aggregate Ethernet interfaces, thereby enabling all the interfaces that you need for your configuration in one step.

Configuration

Step-by-Step Procedure

To configure the number of aggregated Ethernet interfaces on a device:

1. Set the number of aggregated Ethernet interfaces.


```
[edit]
user@host# set chassis aggregated-devices ethernet device-count 2
```
2. If you are done configuring the device, commit the configuration.


```
[edit]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show chassis aggregated-devices** command.

Related Documentation

- [Understanding the Aggregated Ethernet Interfaces Device Count on page 136](#)
- [Aggregated Ethernet Interfaces Configuration Overview on page 136](#)
- [Example: Deleting Aggregated Ethernet Interfaces on page 143](#)
- [Verifying Aggregated Ethernet Interfaces on page 145](#)

Understanding Physical Interfaces for Aggregated Ethernet Interfaces

Supported Platforms [SRX Series](#)

You associate a physical interface with an aggregated Ethernet interface. Doing so associates the physical child links with the logical aggregated parent interface to form a link aggregation group (LAG). You must also specify the constituent physical links by including the **802.3ad** configuration statement.

A physical interface can be added to any aggregated Ethernet interface as long as all member links have the same link speed and the maximum number of member links does not exceed 16. The aggregated Ethernet interface instance number *aex* can be from 0 through 127, for a total of 128 aggregated interfaces.



NOTE: If you specify (on purpose or accidentally) that a link already associated with an aggregated Ethernet interface be associated with another aggregated Ethernet interface, the link is removed from the previous interface (there is no need for you to explicitly delete it) and it is added to the other one.

Related Documentation

- [Understanding Aggregated Ethernet Interfaces on page 133](#)
- [Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces on page 138](#)

Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces

Supported Platforms [SRX Series](#)

This example shows how to associate physical interfaces with aggregated Ethernet interfaces.

- [Requirements on page 139](#)
- [Overview on page 139](#)

- [Configuration on page 139](#)
- [Verification on page 139](#)

Requirements

Before you begin, set the number of aggregated Ethernet interfaces on the device. See [“Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device” on page 137](#).

Overview

In this example, you associate the physical child link of the ge-1/0/0 and ge-2/0/0 physical interfaces with the logical aggregate parent, ae0, thereby creating a LAG. Similarly, you create a LAG that associate the ge-3/0/0, ge-3/0/1, and ge-4/0/1 physical interfaces with the ae1 aggregated Ethernet interface.

Configuration

Step-by-Step Procedure

To associate physical interfaces with aggregated Ethernet interfaces:

1. Create the first LAG.

```
[edit]
user@host# set interfaces ge-1/0/0 gigether-options 802.3ad ae0
user@host# set interfaces ge-2/0/0 gigether-options 802.3ad ae0
```

2. Create the second LAG.

```
[edit]
user@host# set interfaces ge-3/0/0 gigether-options 802.3ad ae1
user@host# set interfaces ge-3/0/1 gigether-options 802.3ad ae1
user@host# sset interfaces ge-4/0/0 gigether-options 802.3ad ae1
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show interfaces** command.

Related Documentation

- [Understanding Physical Interfaces for Aggregated Ethernet Interfaces on page 138](#)
- [Aggregated Ethernet Interfaces Configuration Overview on page 136](#)
- [Verifying Aggregated Ethernet Interfaces on page 145](#)

Understanding Aggregated Ethernet Interface Link Speed

Supported Platforms **SRX Series**

On aggregated Ethernet interfaces, you can set the required link speed for all interfaces included in the bundle. All interfaces that make up a bundle must be the same speed. If you include in the aggregated Ethernet interface an individual link that has a speed

different from the speed you specify in the **link-speed** parameter, an error message will be logged.

The speed value is specified in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).

Aggregated Ethernet interfaces on SRX3000 and SRX5000 line devices can have one of the following speed values:

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

Related Documentation

- [Understanding Aggregated Ethernet Interfaces on page 133](#)
- [Example: Configuring Aggregated Ethernet Link Speed on page 140](#)
- [Understanding Minimum Links for Aggregated Ethernet Interfaces on page 141](#)

Example: Configuring Aggregated Ethernet Link Speed

Supported Platforms [SRX Series](#)

This example shows how to configure the aggregated Ethernet link speed.

- [Requirements on page 140](#)
- [Overview on page 140](#)
- [Configuration on page 140](#)
- [Verification on page 141](#)

Requirements

Before you begin:

- Add the aggregated Ethernet interfaces using the device count. See [“Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device” on page 137](#).
- Associate physical interfaces with the aggregated Ethernet Interfaces. See [“Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces” on page 138](#).

Overview

In this example, you set the required link speed for all interfaces included in the bundle to 10 Gbps. All interfaces that make up a bundle must be the same speed.

Configuration

Step-by-Step Procedure

To configure the aggregated Ethernet link speed:

1. Set the link speed.

```
[edit]
user@host# set interfaces ae0 aggregated-ether-options link-speed 10g
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show interfaces** command.

Related Documentation

- [Understanding Aggregated Ethernet Interface Link Speed on page 139](#)
- [Aggregated Ethernet Interfaces Configuration Overview on page 136](#)
- [Verifying Aggregated Ethernet Interfaces on page 145](#)

Understanding Minimum Links for Aggregated Ethernet Interfaces

Supported Platforms [SRX Series](#)

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

On SRX3000 and SRX5000 line devices, the valid range for the minimum links number is 1 through 16. When the maximum value (16) is specified, all configured links of a bundle must be up for the bundle to be labeled as up.

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

Related Documentation

- [Understanding Aggregated Ethernet Interfaces on page 133](#)
- [Example: Configuring Aggregated Ethernet Minimum Links on page 141](#)
- [Understanding Aggregated Ethernet Interface Link Speed on page 139](#)

Example: Configuring Aggregated Ethernet Minimum Links

Supported Platforms [SRX Series](#)

This example shows how to configure the minimum number of links on an aggregated Ethernet interface that must be up for the bundle as a whole to be labeled as up.

- [Requirements on page 142](#)
- [Overview on page 142](#)
- [Configuration on page 142](#)
- [Verification on page 142](#)

Requirements

Before you begin:

- Add the aggregated Ethernet interfaces using the device count. See [“Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device”](#) on page 137.
- Associate physical interfaces with the aggregated Ethernet Interfaces. See [“Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces”](#) on page 138.
- Configure the aggregated Ethernet link speed. See [“Example: Configuring Aggregated Ethernet Link Speed”](#) on page 140.

Overview

In this example, you specify that on interface ae0 at least eight links must be up for the bundle as a whole to be labeled as up.

Configuration

Step-by-Step Procedure

To configure the minimum number of links on an aggregated Ethernet interface:

1. Set the minimum number of links.

[edit]

```
user@host# set interfaces ae0 aggregated-ether-options minimum-links 8
```

2. If you are done configuring the device, commit the configuration.

[edit]

```
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show interfaces** command.

Related Documentation

- [Understanding Aggregated Ethernet Interface Link Speed](#) on page 139
- [Aggregated Ethernet Interfaces Configuration Overview](#) on page 136
- [Verifying Aggregated Ethernet Interfaces](#) on page 145

Understanding Aggregated Ethernet Interface Removal

Supported Platforms [SRX Series](#)

You can delete an aggregated Ethernet interface from the interface configuration. Junos OS removes the configuration statements related to **aex** and sets this interface to the down state. The deleted aggregated Ethernet interface still exists, but it becomes an empty interface.

Related Documentation

- [Understanding Aggregated Ethernet Interfaces](#) on page 133
- [Example: Deleting Aggregated Ethernet Interfaces](#) on page 143

- [Example: Deleting Aggregated Ethernet Interface Contents on page 144](#)

Example: Deleting Aggregated Ethernet Interfaces

Supported Platforms [SRX Series](#)

This example shows how to delete aggregated Ethernet interfaces using the device count.

- [Requirements on page 143](#)
- [Overview on page 143](#)
- [Configuration on page 143](#)
- [Verification on page 143](#)

Requirements

Before you begin, set the number of aggregated Ethernet interfaces on the device. See “[Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device](#)” on page 137.

Overview

This example shows how to clean up unused aggregated Ethernet interfaces. In this example, you reduce the number of interfaces from 10 to 6, thereby removing the last 4 interfaces from the interface object list.

Configuration

Step-by-Step Procedure

To delete an interface:

1. Set the number of aggregated Ethernet interfaces.


```
[edit]
user@host# delete chassis aggregated-devices ethernet device-count 6
```
2. If you are done configuring the device, commit the configuration.


```
[edit]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show chassis aggregated-devices** command.

Related Documentation

- [Aggregated Ethernet Interfaces Configuration Overview on page 136](#)
- [Example: Deleting Aggregated Ethernet Interface Contents on page 144](#)
- [Verifying Aggregated Ethernet Interfaces on page 145](#)

Example: Deleting Aggregated Ethernet Interface Contents

Supported Platforms [SRX Series](#)

This example shows how to delete the contents of an aggregated Ethernet interface.

- [Requirements on page 144](#)
- [Overview on page 144](#)
- [Configuration on page 144](#)
- [Verification on page 144](#)

Requirements

Before you begin:

- Set the number of aggregated Ethernet interfaces on the device. See [“Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device” on page 137](#).
- Associate a physical interface with the aggregated Ethernet interface. See [“Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces” on page 138](#).
- Set the required link speed for all the interfaces included in the bundle. See [“Example: Configuring Aggregated Ethernet Link Speed” on page 140](#).
- Configure the minimum number of links that must be up for the bundle as a whole to be labeled as up. See [“Example: Configuring Aggregated Ethernet Minimum Links” on page 141](#).

Overview

In this example, you delete the contents of the ae4 aggregated Ethernet interface, which sets it to the down state.

Configuration

Step-by-Step Procedure

To delete the contents of an aggregated Ethernet interface:

1. Delete the interface.

```
[edit]  
user@host# delete interfaces ae4
```
2. If you are done configuring the device, commit the configuration.

```
[edit]  
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show interfaces** command.

Related Documentation

- [Aggregated Ethernet Interfaces Configuration Overview on page 136](#)
- [Example: Deleting Aggregated Ethernet Interfaces on page 143](#)

- [Verifying Aggregated Ethernet Interfaces on page 145](#)

Verifying Aggregated Ethernet Interfaces

Supported Platforms [SRX Series](#)

- [Verifying Aggregated Ethernet Interfaces \(terse\) on page 145](#)
- [Verifying Aggregated Ethernet Interfaces \(extensive\) on page 145](#)

Verifying Aggregated Ethernet Interfaces (terse)

Supported Platforms [SRX Series](#)

Purpose Display status information in terse (concise) format for aggregated Ethernet interfaces.

Action From operational mode, enter the **show interfaces ae0 terse** command.

```
user@host> show interfaces ae0 terse
ge-2/0/0.0          up    up    aenet  --> ae0.0
ge-2/0/0.32767     up    up    aenet  --> ae0.32767
ge-2/0/1.0         up    up    aenet  --> ae0.0
ge-2/0/1.32767     up    up    aenet  --> ae0.32767
ae0                up    up
ae0.0              up    up    bridge
ae0.32767          up    up    multiservice
```

The output shows the bundle relationship for the aggregated Ethernet interface and the overall status of the interface, including the following information:

- The link aggregation control PDUs run on the .0 child logical interfaces for the untagged aggregated Ethernet interface.
- The link aggregation control PDUs run on the .32767 child logical interfaces for the VLAN-tagged aggregated Ethernet interface.
- The .32767 logical interface is created for the parent link and all child links.

Verifying Aggregated Ethernet Interfaces (extensive)

Supported Platforms [SRX Series](#)

Purpose Display status information and statistics in extensive (detailed) format for aggregated Ethernet interfaces.

Action From operational mode, enter the **show interfaces ae0 extensive** command.

```
user@host> show interfaces ae0 extensive
Physical interface: ae0, Enabled, Physical link is Up
...
Logical interface ae0.0 (Index 67) (SNMP ifIndex 628) (Generation 134)
...
LACP info:          Role    System          System    Port    Port    Port
                   priority identifier priority number  key
```

```

ge-5/0/0.0 Actor 127 00:1f:12:8c:af:c0 127 832 1
ge-5/0/0.0 Partner 127 00:1f:12:8f:d7:c0 127 640 1
ge-5/0/1.0 Actor 127 00:1f:12:8c:af:c0 127 833 1
ge-5/0/1.0 Partner 127 00:1f:12:8f:d7:c0 127 641 1

LACP Statistics: LACP Rx LACP Tx Unknown Rx Illegal Rx
ge-5/0/0.0 12830 7090 0 0
ge-5/0/1.0 10304 4786 0 0
...
Logical interface ae0.32767 (Index 70) (SNMP ifIndex 630) (Generation 135)
...
LACP info: Role System System Port Port Port
           priority identifier priority number key

ge-5/0/0.32767 Actor 127 00:1f:12:8c:af:c0 127 832 1
ge-5/0/0.32767 Partner 127 00:1f:12:8f:d7:c0 127 640 1
ge-5/0/1.32767 Actor 127 00:1f:12:8c:af:c0 127 833 1
ge-5/0/1.32767 Partner 127 00:1f:12:8f:d7:c0 127 641 1

LACP Statistics: LACP Rx LACP Tx Unknown Rx Illegal Rx
ge-5/0/0.32767 12830 7090 0 0
ge-5/0/1.32767 10304 4786 0 0
...

```

The output shows detailed aggregated Ethernet interface information. This portion of the output shows LACP information and LACP statistics for each logical aggregated Ethernet interface.

Related Documentation • [Aggregated Ethernet Interfaces Configuration Overview on page 136](#)

Understanding VLAN Tagging for Aggregated Ethernet Interfaces

Supported Platforms SRX1500, SRX5400, SRX5600, SRX5800

Aggregated Ethernet interfaces can be either VLAN-tagged or untagged, with LACP enabled or disabled. Aggregated Ethernet interfaces on the SRX3000 and SRX5000 lines support the configuration of **native-vlan-id**, which consists of the following configuration statements:

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

- `swap-push`
- `swap-swap`

**Related
Documentation**

- [Understanding Aggregated Ethernet Interfaces on page 133](#)
- [Aggregated Ethernet Interfaces Configuration Overview on page 136](#)

Understanding Promiscuous Mode for Aggregated Ethernet Interfaces

Supported Platforms [SRX1500, SRX5400, SRX5600, SRX5800](#)

You can enable promiscuous mode on aggregated Ethernet interfaces. When promiscuous mode is enabled on a Layer 3 Ethernet interface, all packets received on the interface are sent to the central point or Services Processing Unit (SPU) regardless of the destination MAC address of the packet. If you enable promiscuous mode on an aggregated Ethernet interface, promiscuous mode is then enabled on all member interfaces.

**Related
Documentation**

- [Understanding Aggregated Ethernet Interfaces on page 133](#)
- [Aggregated Ethernet Interfaces Configuration Overview on page 136](#)

Configuring Link Aggregation Control Protocol

- [Understanding LACP on Standalone Devices on page 149](#)
- [Example: Configuring LACP on Standalone Devices on page 150](#)
- [Verifying LACP on Standalone Devices on page 151](#)
- [Understanding LACP on Chassis Clusters on page 153](#)
- [Example: Configuring LACP on Chassis Clusters on page 155](#)
- [Verifying LACP on Redundant Ethernet Interfaces on page 157](#)
- [LAG and LACP Support on SRX5000 Line Devices with I/O Cards \(IOCs\) on page 158](#)
- [Example: Configuring LAG Interface on an SRX5000 Line Device with IOC2 or IOC3 on page 160](#)

Understanding LACP on Standalone Devices

Supported Platforms [SRX Series](#)

Link Aggregation Control Protocol (LACP) provides a standardized means for exchanging information between partner systems on a link. Within LACP, the local end of a child link is known as the actor and the remote end of the link is known as the partner.

LACP is enabled on an aggregated Ethernet interface by setting the mode to either passive or active. However, to initiate the transmission of link aggregation control protocol data units (PDUs) and response link aggregation control PDUs, you must enable LACP at both the local and remote ends of the links, and one end must be active:

- **Active mode**—If either the actor or partner is active, they exchange link aggregation control PDUs. The actor sends link aggregation control PDUs to its protocol partner that convey what the actor knows about its own state and that of the partner's state.
- **Passive mode**—If the actor and partner are both in passive mode, they do not exchange link aggregation control PDUs. As a result, the aggregated Ethernet links do not come up. In passive transmission mode, links send out link aggregation control PDUs only when they receive them from the remote end of the same link.

By default, the actor and partner transmit link aggregation control PDUs every second. You can configure different periodic rates on active and passive interfaces. When you

configure the active and passive interfaces at different rates, the transmitter honors the receiver's rate.

You configure the interval at which the interfaces on the remote side of the link transmit link aggregation control PDUs by configuring the **periodic** statement on the interfaces on the local side. It is the configuration on the local side that specifies the behavior of the remote side. That is, the remote side transmits link aggregation control PDUs at the specified interval. The interval can be **fast** (every second) or **slow** (every 30 seconds).



NOTE: On all high-end SRX Series devices, the LACP is not supported on Layer 2 interfaces.

Related Documentation

- [Understanding Aggregated Ethernet Interfaces on page 133](#)
- [Understanding LACP on Chassis Clusters on page 153](#)
- [Example: Configuring LACP on Standalone Devices on page 150](#)

Example: Configuring LACP on Standalone Devices

Supported Platforms [SRX Series](#)

This example shows how to configure LACP on standalone devices.

- [Requirements on page 150](#)
- [Overview on page 150](#)
- [Configuration on page 151](#)
- [Verification on page 151](#)

Requirements

Before you begin:

- Add the aggregated Ethernet interfaces using the device count. See “[Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device](#)” on page 137.
- Associate physical interfaces with the aggregated Ethernet Interfaces. See “[Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces](#)” on page 138.
- Configure the aggregated Ethernet link speed. See “[Example: Configuring Aggregated Ethernet Link Speed](#)” on page 140.
- Configure the aggregated Ethernet minimum links speed. See “[Example: Configuring Aggregated Ethernet Minimum Links](#)” on page 141.

Overview

In this example, you set LACP to passive mode for the ae0 interface. You set the LACP mode for the ae1 interface to active and set the link aggregation control PDU transmit interval to slow, which is every 30 seconds.

Configuration

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure LACP on standalone devices:

1. Set the first LACP.


```
[edit interfaces]
user@host# set ae0 aggregated-ether-options lACP passive
```
2. Set the second LACP.


```
[edit interfaces]
user@host# set ae1 aggregated-ether-options lACP active
user@host# set ae1 aggregated-ether-options lACP periodic slow
```
3. If you are done configuring the device, commit the configuration.


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

Verification

To confirm that the configuration is working properly, perform these tasks:

- Related Documentation**
- [Understanding LACP on Standalone Devices on page 149](#)
 - [Verifying LACP on Standalone Devices on page 151](#)

Verifying LACP on Standalone Devices

Supported Platforms [SRX Series](#)

- [Verifying LACP Statistics on page 151](#)
- [Verifying LACP Aggregated Ethernet Interfaces on page 152](#)

Verifying LACP Statistics

Supported Platforms [SRX Series](#)

Purpose Display LACP statistics for aggregated Ethernet interfaces.

Action From operational mode, enter the `show lACP statistics interfaces ae0` command.

```
user@host> show lACP statistics interfaces ae0
Aggregated interface: ae0
LACP Statistics:      LACP Rx      LACP Tx      Unknown Rx      Illegal Rx
ge-2/0/0              1352         2035          0                0
ge-2/0/1              1352         2056          0                0
ge-2/2/0              1352         2045          0                0
ge-2/2/1              1352         2043          0                0
```

The output shows LACP statistics for each physical interface associated with the aggregated Ethernet interface, such as the following:

- The LACP received counter that increments for each normal hello
- The number of LACP transmit packet errors logged
- The number of unrecognized packet errors logged
- The number of invalid packets received

Use the following command to clear the statistics and see only new changes:

```
user@host# clear lACP statistics interfaces ae0
```

Verifying LACP Aggregated Ethernet Interfaces

Supported Platforms [SRX Series](#)

Purpose Display LACP status information for aggregated Ethernet interfaces.

Action From operational mode, enter the **show lACP interfaces ae0** command.

```
user@host> show lACP interfaces ae0
Aggregated interface: ae0
  LACP state:      Role   Exp   Def   Dist  Col   Syn   Aggr  Timeout  Activity
  ge-2/0/0        Actor  No    No    Yes   Yes   Yes   Yes    Fast    Active
  ge-2/0/0        Partner No    No    Yes   Yes   Yes   Yes    Fast    Active
  ge-2/0/1        Actor  No    No    Yes   Yes   Yes   Yes    Fast    Active
  ge-2/0/1        Partner No    No    Yes   Yes   Yes   Yes    Fast    Active
  ge-2/2/0        Actor  No    No    Yes   Yes   Yes   Yes    Fast    Active
  ge-2/2/0        Partner No    No    Yes   Yes   Yes   Yes    Fast    Active
  ge-2/2/1        Actor  No    No    Yes   Yes   Yes   Yes    Fast    Active
  ge-2/2/1        Partner No    No    Yes   Yes   Yes   Yes    Fast    Active
  LACP protocol:  Receive State  Transmit State  Mux State
  ge-2/0/0        Current    Fast periodic  Collecting distributing
  ge-2/0/1        Current    Fast periodic  Collecting distributing
  ge-2/2/0        Current    Fast periodic  Collecting distributing
  ge-2/2/1        Current    Fast periodic  Collecting distributing
```

The output shows aggregated Ethernet interface information, including the following information:

- The LACP state—Indicates whether the link in the bundle is an actor (local or near-end of the link) or a partner (remote or far-end of the link).
- The LACP mode—Indicates whether both ends of the aggregated Ethernet interface are enabled (active or passive)—at least one end of the bundle must be active.
- The periodic link aggregation control PDU transmit rate.
- The LACP protocol state—Indicates the link is up if it is collecting and distributing packets.

- Related Documentation**
- [Example: Configuring LACP on Standalone Devices on page 150](#)
 - [Verifying LACP on Redundant Ethernet Interfaces on page 157](#)

Understanding LACP on Chassis Clusters

Supported Platforms [SRX Series](#)

You can combine multiple physical Ethernet ports to form a logical point-to-point link, known as a link aggregation group (LAG) or bundle, such that a media access control (MAC) client can treat the LAG as if it were a single link.

LAGs can be established across nodes in a chassis cluster to provide increased interface bandwidth and link availability.

The Link Aggregation Control Protocol (LACP) provides additional functionality for LAGs. LACP is supported in standalone deployments, where aggregated Ethernet interfaces are supported, and in chassis cluster deployments, where aggregated Ethernet interfaces and redundant Ethernet interfaces are supported simultaneously.

You configure LACP on a redundant Ethernet interface by setting the LACP mode for the parent link with the `lACP` statement. The LACP mode can be off (the default), active, or passive.

This topic contains the following sections:

- [Chassis Cluster Redundant Ethernet Interface Link Aggregation Groups on page 153](#)
- [Sub-LAGs on page 154](#)
- [Supporting Hitless Failover on page 154](#)
- [Managing Link Aggregation Control PDUs on page 155](#)

Chassis Cluster Redundant Ethernet Interface Link Aggregation Groups

A redundant Ethernet interface has active and standby links located on two nodes in a chassis cluster. All active links are located on one node, and all standby links are located on the other node. You can configure up to eight active links and eight standby links per node.

When at least two physical child interface links from each node are included in a redundant Ethernet interface configuration, the interfaces are combined within the redundant Ethernet interface to form a redundant Ethernet interface LAG.

Having multiple active redundant Ethernet interface links reduces the possibility of failover. For example, when an active link is out of service, all traffic on this link is distributed to other active redundant Ethernet interface links, instead of triggering a redundant Ethernet active/standby failover.

Aggregated Ethernet interfaces, known as local LAGs, are also supported on either node of a chassis cluster but cannot be added to redundant Ethernet interfaces. Likewise, any child interface of an existing local LAG cannot be added to a redundant Ethernet interface, and vice versa. The total maximum number of combined individual node LAG interfaces (ae) and redundant Ethernet (reth) interfaces per cluster is 128.

However, aggregated Ethernet interfaces and redundant Ethernet interfaces can coexist, because the functionality of a redundant Ethernet interface relies on the Junos OS aggregated Ethernet framework.

For more information, see *Understanding Chassis Cluster Redundant Ethernet Interface Link Aggregation Groups for Branch SRX Series Devices* or *Understanding Chassis Cluster Redundant Ethernet Interface Link Aggregation Groups for High-End SRX Series Devices*.

Minimum Links

Redundant Ethernet interface configuration includes a **minimum-links** setting that allows you to set a minimum number of physical child links in a redundant Ethernet interface LAG that must be working on the primary node for the interface to be up. The default **minimum-links** value is 1. When the number of physical links on the primary node in a redundant Ethernet interface falls below the **minimum-links** value, the interface might be down even if some links are still working. For more information, see *Example: Configuring Chassis Cluster Minimum Links*.

Sub-LAGs

LACP maintains a point-to-point LAG. Any port connected to the third point is denied. However, a redundant Ethernet interface does connect to two different systems or two remote aggregated Ethernet interfaces by design.

To support LACP on both redundant Ethernet interface active and standby links, a redundant Ethernet interface can be modeled to consist of two sub-LAGs, where all active links form an active sub-LAG and all standby links form a standby sub-LAG.

In this model, LACP selection logic is applied and limited to one sub-LAG at a time. In this way, two redundant Ethernet interface sub-LAGs are maintained simultaneously while all the LACP advantages are preserved for each sub-LAG.

It is necessary for the switches used to connect the nodes in the cluster to have a LAG link configured and 802.3ad enabled for each LAG on both nodes so that the aggregate links are recognized as such and correctly pass traffic.



NOTE: The redundant Ethernet interface LAG child links from each node in the chassis cluster must be connected to a different LAG at the peer devices. If a single peer switch is used to terminate the redundant Ethernet interface LAG, two separate LAGs must be used in the switch.

Supporting Hitless Failover

With LACP, the redundant Ethernet interface supports hitless failover between the active and standby links in normal operation. The term *hitless* means that the redundant Ethernet interface state remains up during a failover.

The lacpd process manages both the active and standby links of the redundant Ethernet interfaces. A redundant Ethernet interface state remains up when the number of active up links is more than the number of minimum links configured. Therefore, to support

hitless failover, the LACP state on the redundant Ethernet interface standby links must be collected and distributed before failover occurs.

Managing Link Aggregation Control PDUs

The protocol data units (PDUs) contain information about the state of the link. By default, aggregated and redundant Ethernet links do not exchange link aggregation control PDUs.

You can configure PDUs exchange in the following ways:

- Configure Ethernet links to actively transmit link aggregation control PDUs
- Configure Ethernet links to passively transmit PDUs, sending out link aggregation control PDUs only when they are received from the remote end of the same link

The local end of a child link is known as the actor and the remote end of the link is known as the partner. That is, the actor sends link aggregation control PDUs to its protocol partner that convey what the actor knows about its own state and that of the partner's state.

You configure the interval at which the interfaces on the remote side of the link transmit link aggregation control PDUs by configuring the **periodic** statement on the interfaces on the local side. It is the configuration on the local side that specifies the behavior of the remote side. That is, the remote side transmits link aggregation control PDUs at the specified interval. The interval can be **fast** (every second) or **slow** (every 30 seconds).

For more information, see [“Example: Configuring LACP on Chassis Clusters” on page 155](#).

By default, the actor and partner transmit link aggregation control PDUs every second. You can configure different periodic rates on active and passive interfaces. When you configure the active and passive interfaces at different rates, the transmitter honors the receiver's rate.

Related Documentation

- [Example: Configuring LACP on Chassis Clusters on page 155](#)

Example: Configuring LACP on Chassis Clusters

Supported Platforms [SRX Series](#)

This example shows how to configure LACP on chassis clusters.

- [Requirements on page 155](#)
- [Overview on page 156](#)
- [Configuration on page 156](#)
- [Verification on page 156](#)

Requirements

Before you begin:

- Add the aggregated Ethernet interfaces using the device count. See “[Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device](#)” on page 137.
- Associate physical interfaces with the aggregated Ethernet Interfaces. See “[Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces](#)” on page 138.
- Configure the aggregated Ethernet link speed. See “[Example: Configuring Aggregated Ethernet Link Speed](#)” on page 140.
- Configure the aggregated Ethernet minimum links speed. See “[Example: Configuring Aggregated Ethernet Minimum Links](#)” on page 141.
- Configure the LACP on standalone devices. See “[Example: Configuring LACP on Standalone Devices](#)” on page 150.

Overview

In this example, you set LACP to passive mode for the reth0 interface. You set the LACP mode for the reth1 interface to active and set the link aggregation control PDU transmit interval to slow, which is every 30 seconds.

Configuration

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see the *CLI User Guide*.

To configure LACP on chassis clusters:

1. Set the first LACP on primary node1.

```
[edit interfaces]
user@host# set reth0 redundant-ether-options lacp passive
```

2. Set the second LACP.

```
[edit interfaces]
user@host# set reth1 redundant-ether-options lacp active
user@host# set reth1 redundant-ether-options lacp periodic slow
```

3. If you are done configuring the device, commit the configuration.

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

Verification

Verifying LACP on Redundant Ethernet Interfaces

Purpose Display LACP status information for redundant Ethernet interfaces.

Action From operational mode, enter the **show lacp interfaces reth0** command.

```
user@host> show lacp interfaces reth0
Aggregated interface: reth0
LACP state:      Role  Exp  Def  Dist  Col  Syn  Aggr  Timeout  Activity
ge-11/0/0       Actor No   No   Yes  Yes  Yes  Yes   Fast    Active
ge-11/0/0       Partner No   No   Yes  Yes  Yes  Yes   Fast    Active
```

```

ge-11/0/1    Actor    No    No    Yes    Yes    Yes    Yes    Fast    Active
ge-11/0/1    Partner  No    No    Yes    Yes    Yes    Yes    Fast    Active
ge-11/0/2    Actor    No    No    Yes    Yes    Yes    Yes    Fast    Active
ge-11/0/2    Partner  No    No    Yes    Yes    Yes    Yes    Fast    Active
ge-11/0/3    Actor    No    No    Yes    Yes    Yes    Yes    Fast    Active
ge-11/0/3    Partner  No    No    Yes    Yes    Yes    Yes    Fast    Active
ge-3/0/0     Actor    No    No    Yes    Yes    Yes    Yes    Fast    Active
ge-3/0/0     Partner  No    No    Yes    Yes    Yes    Yes    Fast    Active
ge-3/0/1     Actor    No    No    Yes    Yes    Yes    Yes    Fast    Active
ge-3/0/1     Partner  No    No    Yes    Yes    Yes    Yes    Fast    Active
ge-3/0/2     Actor    No    No    Yes    Yes    Yes    Yes    Fast    Active
ge-3/0/2     Partner  No    No    Yes    Yes    Yes    Yes    Fast    Active
ge-3/0/3     Actor    No    No    Yes    Yes    Yes    Yes    Fast    Active
ge-3/0/3     Partner  No    No    Yes    Yes    Yes    Yes    Fast    Active
LACP protocol:      Receive State  Transmit State  Mux State
ge-11/0/0           Current      Fast periodic  Collecting distributing
ge-11/0/1           Current      Fast periodic  Collecting distributing
ge-11/0/2           Current      Fast periodic  Collecting distributing
ge-11/0/3           Current      Fast periodic  Collecting distributing
ge-3/0/0            Current      Fast periodic  Collecting distributing
ge-3/0/1            Current      Fast periodic  Collecting distributing
ge-3/0/2            Current      Fast periodic  Collecting distributing
ge-3/0/3            Current      Fast periodic  Collecting distributing
{primary:node1}

```

The output shows redundant Ethernet interface information, such as the following:

- The LACP state—Indicates whether the link in the bundle is an actor (local or near-end of the link) or a partner (remote or far-end of the link).
- The LACP mode—Indicates whether both ends of the aggregated Ethernet interface are enabled (active or passive)—at least one end of the bundle must be active.
- The periodic link aggregation control PDU transmit rate.
- The LACP protocol state—Indicates the link is up if it is collecting and distributing packets.

Related Documentation

- [Understanding LACP on Chassis Clusters on page 153](#)
- [Verifying LACP on Redundant Ethernet Interfaces on page 157](#)

Verifying LACP on Redundant Ethernet Interfaces

Supported Platforms [SRX Series](#)

Purpose Display LACP status information for redundant Ethernet interfaces.

Action From operational mode, enter the **show lacp interfaces reth0** command.

```

user@host> show lacp interfaces reth0
Aggregated interface: reth0
LACP state:
ge-11/0/0     Actor    No    No    Yes    Yes    Yes    Yes    Fast    Active
ge-11/0/0     Partner  No    No    Yes    Yes    Yes    Yes    Fast    Active
ge-11/0/1     Actor    No    No    Yes    Yes    Yes    Yes    Fast    Active
ge-11/0/1     Partner  No    No    Yes    Yes    Yes    Yes    Fast    Active
ge-11/0/2     Actor    No    No    Yes    Yes    Yes    Yes    Fast    Active

```

ge-11/0/2	Partner	No	No	Yes	Yes	Yes	Yes	Fast	Active
ge-11/0/3	Actor	No	No	Yes	Yes	Yes	Yes	Fast	Active
ge-11/0/3	Partner	No	No	Yes	Yes	Yes	Yes	Fast	Active
ge-3/0/0	Actor	No	No	Yes	Yes	Yes	Yes	Fast	Active
ge-3/0/0	Partner	No	No	Yes	Yes	Yes	Yes	Fast	Active
ge-3/0/1	Actor	No	No	Yes	Yes	Yes	Yes	Fast	Active
ge-3/0/1	Partner	No	No	Yes	Yes	Yes	Yes	Fast	Active
ge-3/0/2	Actor	No	No	Yes	Yes	Yes	Yes	Fast	Active
ge-3/0/2	Partner	No	No	Yes	Yes	Yes	Yes	Fast	Active
ge-3/0/3	Actor	No	No	Yes	Yes	Yes	Yes	Fast	Active
ge-3/0/3	Partner	No	No	Yes	Yes	Yes	Yes	Fast	Active
LACP protocol:		Receive State		Transmit State				Mux State	
ge-11/0/0		Current		Fast periodic		Collecting		distributing	
ge-11/0/1		Current		Fast periodic		Collecting		distributing	
ge-11/0/2		Current		Fast periodic		Collecting		distributing	
ge-11/0/3		Current		Fast periodic		Collecting		distributing	
ge-3/0/0		Current		Fast periodic		Collecting		distributing	
ge-3/0/1		Current		Fast periodic		Collecting		distributing	
ge-3/0/2		Current		Fast periodic		Collecting		distributing	
ge-3/0/3		Current		Fast periodic		Collecting		distributing	

{primary:node1}

The output shows redundant Ethernet interface information, such as the following:

- The LACP state—Indicates whether the link in the bundle is an actor (local or near-end of the link) or a partner (remote or far-end of the link).
- The LACP mode—Indicates whether both ends of the aggregated Ethernet interface are enabled (active or passive)—at least one end of the bundle must be active.
- The periodic link aggregation control PDU transmit rate.
- The LACP protocol state—Indicates the link is up if it is collecting and distributing packets.

Related Documentation

- [Example: Configuring LACP on Chassis Clusters on page 155](#)
- [Verifying LACP on Standalone Devices on page 151](#)

LAG and LACP Support on SRX5000 Line Devices with I/O Cards (IOCs)

Supported Platforms [SRX5400, SRX5600, SRX5800](#)

LAG and LACP Support on the SRX5000 Module Port Concentrator

The SRX5000 Module Port Concentrator (SRX5K-MPC) on SRX5400, SRX5600, and SRX5800 devices supports link aggregation groups (LAGs) and Link Aggregation Control Protocol (LACP).

Support for LAGs based on IEEE 802.3ad makes it possible to aggregate physical interface links on your device. LAGs provide increased interface bandwidth and link availability by linking physical ports and load-balancing traffic crossing the combined interface.

LACP provides a standardized means for exchanging information between partner (remote or far-end of the link) systems on a link. This exchange allows their link aggregation control instances to reach agreement on the identity of the LAG to which the link belongs,

and then to move the link to that LAG. This exchange also enables the transmission and reception processes for the link to function in an orderly manner.

The following LAG and LACP features are supported on the SRX5K-MPC:

- Bandwidth aggregation—Increases bandwidth, provides graceful degradation as failure occurs, and increases availability.
- Link redundancy and load balancing (within chassis cluster)—Provides network redundancy by load-balancing traffic across all available links. If one of the links should fail, the system automatically load-balances traffic across all remaining links.
- Dynamic link management—Enables automatic addition and deletion of individual links to the aggregate bundle without user intervention.

LACP supports the following features:

- LACP bundles several physical interfaces to form one logical interface by exchanging LACP packets between the local interface and the remote interface. LACP monitors the link for changes in interface state by exchanging a periodic LACP heartbeat between two sides. Any changes in interface state are reflected in the LACP packet.
- Normally after an LACP is configured and committed, two sides start to exchange interface and port information. Once they identify each other and match the LACP state machine criteria, the LACP is declared as up. You can deactivate or delete the LACP configuration.
- By default, the LACP packets are exchanged in every second. You can configure the LACP interval as fast (every second) or slow (every 30 seconds) to ensure the health of the interfaces.
- LACP supports distributed and centralized modes. Chassis cluster setup is recommended to operate with LACP distributed mode, which handles chassis cluster failover better. The centralized mode might experience traffic loss during failover.

SRX5K-MPCs on SRX5000 line devices provide active and standby support with redundant Ethernet interface LAGs in chassis cluster deployments.

LAG and LACP Support on the SRX5000 Line IOCs in Express Path Mode

The IOC2 and IOC3 cards on SRX5400, SRX5600, and SRX5800 devices support link aggregation groups (LAGs) and Link Aggregation Control Protocol (LACP) in Express path mode.

You can use the links in a LAG as ingress or egress interfaces in Express Path mode. The LAG links can include links from cards such as IOC2 or IOC3. For a LAG link to qualify for Express Path, all its member links should be connected to Express Path-enabled network processors. If Express Path is disabled on any of the member links in a LAG, a regular session (non-Express Path session) is created.

**NOTE:**

- Cross-IOC LAG interfaces do not support Layer 2 transparent mode.
- Mixed interface speeds are not supported on the same aggregated bundle.
- A redundant Ethernet interface or aggregated Ethernet interface must contain child interfaces from the same IOC type. For example, if one child link is from 10-Gigabit Ethernet on IOC2, the second child link should also be from IOC2. Similarly, both child interfaces can be from IOC3. Configuring child interfaces by mixing links from both IOC2 and IOC3 is not supported.

Related Documentation

- [Aggregated Ethernet Interfaces Configuration Overview on page 136](#)
- [Example: Configuring LACP on Standalone Devices on page 150](#)
- [Example: Configuring LACP on Chassis Clusters on page 155](#)

Example: Configuring LAG Interface on an SRX5000 Line Device with IOC2 or IOC3

Supported Platforms SRX5400, SRX5600, SRX5800

IEEE 802.3ad link aggregation enables you to group Ethernet interfaces to form a single, aggregated Ethernet interface, also known as a LAG or bundle. The LACP provides additional functionality for LAGs.

This example shows how to configure LAG on an SRX Series device using the links from either IOC2 or IOC3 in Express Path mode.

- [Requirements on page 160](#)
- [Overview on page 160](#)
- [Configuration on page 161](#)
- [Verification on page 163](#)

Requirements

This example uses the following software and hardware components:

- Junos OS Release 15.1X49-D40 or later for SRX Series devices.
- SRX5800 with IOC2 or IOC3 with Express Path enabled on IOC2 and IOC3. For details, see *Example: Configuring SRX5K-MPC3-100G10G (IOC3) and SRX5K-MPC3-40G10G (IOC3) on an SRX5000 Line Device to Support Express Path*.

Overview

In this example, you create a logical aggregated Ethernet interface and define the parameters associated with the logical aggregated Ethernet interface, such as a logical unit, interface properties, and LACP. Next, define the member links to be contained within the aggregated Ethernet interface—for example, four 10-Gigabit Ethernet interfaces. Finally, configure an LACP for link detection.

The following member links are used in this example:

- xe-0/0/8
- xe-0/0/9
- xe-1/0/8
- xe-1/0/9
- xe-3/1/4
- xe-3/1/5
- xe-5/1/4
- xe-5/1/5

Configuration

CLI Quick Configuration

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

```
set chassis aggregated-devices ethernet device-count 5
set interfaces xe-0/0/8 gigether-options 802.3ad ae1
set interfaces xe-0/0/9 gigether-options 802.3ad ae0
set interfaces xe-1/0/8 gigether-options 802.3ad ae1
set interfaces xe-1/0/9 gigether-options 802.3ad ae0
set interfaces xe-3/1/4 gigether-options 802.3ad ae1
set interfaces xe-3/1/5 gigether-options 802.3ad ae0
set interfaces xe-5/1/4 gigether-options 802.3ad ae1
set interfaces xe-5/1/5 gigether-options 802.3ad ae0
set interfaces ae0 unit 0 family inet address 17.0.0.1/24
set interfaces ae1 unit 0 family inet address 16.0.0.1/24
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp active
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see the [Junos OS CLI User Guide](#).

To configure LAG Interfaces:

1. Specify the number of aggregated Ethernet interfaces to be created.


```
[edit chassis]
user@host# set aggregated-devices ethernet device-count 5
```
2. Specify the members to be included within the aggregated Ethernet bundle.


```
[edit interfaces]
user@host# set xe-0/0/8 gigether-options 802.3ad ae1
user@host# set xe-0/0/9 gigether-options 802.3ad ae0
user@host# set xe-1/0/8 gigether-options 802.3ad ae1
user@host# set xe-1/0/9 gigether-options 802.3ad ae0
user@host# set xe-3/1/4 gigether-options 802.3ad ae1
user@host# set xe-3/1/5 gigether-options 802.3ad ae0
```

```

user@host# set xe-5/1/4 ggether-options 802.3ad ae1
user@host# set xe-5/1/5 ggether-options 802.3ad ae0

```

3. Assign an IP address to ae0 and ae1.

```

[edit interfaces]
user@host# set ae0 unit 0 family inet address 17.0.0.1/24
user@host# set ae1 unit 0 family inet address 16.0.0.1/24

```

4. Set the LACP on reth0.

```

[edit interfaces]
user@host# set ae0 aggregated-ether-options lacp active
user@host# set ae1 aggregated-ether-options lacp active

```

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

```

[edit]
user@host# show interfaces
xe-0/0/8 {
  ggether-options {
    802.3ad ae1;
  }
}
xe-0/0/9 {
  ggether-options {
    802.3ad ae0;
  }
}
xe-1/0/8 {
  ggether-options {
    802.3ad ae1;
  }
}
xe-1/0/9 {
  ggether-options {
    802.3ad ae0;
  }
}
xe-3/1/4 {
  ggether-options {
    802.3ad ae1;
  }
}
xe-3/1/5 {
  ggether-options {
    802.3ad ae0;
  }
}
ae0 {
  aggregated-ether-options {
    lacp {
      active;
    }
  }
}

```

```

}
unit 0 {
  family inet {
    address 17.0.0.1/24;
  }
}
}
ae1 {
  aggregated-ether-options {
    lacp {
      active;
    }
  }
  unit 0 {
    family inet {
      address 16.0.0.1/24;
    }
  }
}
}

[edit]
user@host# show chassis
aggregated-devices {
  ethernet {
    device-count 5;
  }
}
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying LACP on Redundant Ethernet Interfaces

Purpose Display LACP status information for redundant Ethernet interfaces.

Action From operational mode, enter the **show lacp interfaces** command to check that LACP has been enabled as active on one end.

```

user@host> show lacp interfaces
Aggregated interface: ae0
LACP state:
  Role   Exp  Def  Dist  Co1  Syn  Aggr  Timeout  Activity
xe-0/0/9 Actor No  No   Yes  Yes  Yes  Yes  Fast  Active
xe-0/0/9 Partner No  No   Yes  Yes  Yes  Yes  Fast  Active
xe-1/0/9 Actor No  No   Yes  Yes  Yes  Yes  Fast  Active
xe-1/0/9 Partner No  No   Yes  Yes  Yes  Yes  Fast  Active
xe-3/1/5 Actor No  No   Yes  Yes  Yes  Yes  Fast  Active
xe-3/1/5 Partner No  No   Yes  Yes  Yes  Yes  Fast  Active
xe-5/1/5 Actor No  No   Yes  Yes  Yes  Yes  Fast  Active
xe-5/1/5 Partner No  No   Yes  Yes  Yes  Yes  Fast  Active
LACP protocol:
  Receive State  Transmit State  Mux State
xe-0/0/9        Current        Fast periodic  Collecting distributing
xe-1/0/9        Current        Fast periodic  Collecting distributing
xe-3/1/5        Current        Fast periodic  Collecting distributing
xe-5/1/5        Current        Fast periodic  Collecting distributing

Aggregated interface: ae1

```

```

LACP state:      Role  Exp  Def  Dist  Co1  Syn  Aggr  Timeout  Activity
xe-0/0/8        Actor No   No   Yes  Yes  Yes  Yes    Fast    Active
xe-0/0/8        Partner No  No   Yes  Yes  Yes  Yes    Fast    Active
xe-1/0/8        Actor No   No   Yes  Yes  Yes  Yes    Fast    Active
xe-1/0/8        Partner No  No   Yes  Yes  Yes  Yes    Fast    Active
xe-3/1/4        Actor No   No   Yes  Yes  Yes  Yes    Fast    Active
xe-3/1/4        Partner No  No   Yes  Yes  Yes  Yes    Fast    Active
xe-5/1/4        Actor No   No   Yes  Yes  Yes  Yes    Fast    Active
xe-5/1/4        Partner No  No   Yes  Yes  Yes  Yes    Fast    Active
LACP protocol:  Receive State  Transmit State  Mux State
xe-0/0/8        Current      Fast periodic Collecting distributing
xe-1/0/8        Current      Fast periodic Collecting distributing
xe-3/1/4        Current      Fast periodic Collecting distributing
xe-5/1/4        Current      Fast periodic Collecting distributing
    
```

The output indicates that LACP has been set up correctly and is active at one end.

- Related Documentation**
- [Understanding LACP on Chassis Clusters on page 153](#)
 - [Verifying LACP on Redundant Ethernet Interfaces on page 157](#)

Configuring Gigabit Ethernet Physical Interface Modules

- [Understanding the 1-Port Gigabit Ethernet SFP Mini-PIM on page 165](#)
- [Example: Configuring the 1-Port Gigabit Ethernet SFP Mini-PIM Interface on page 167](#)
- [Understanding the 2-Port 10-Gigabit Ethernet XPIM on page 173](#)
- [Example: Configuring the 2-Port 10-Gigabit Ethernet XPIM Interface on page 176](#)

Understanding the 1-Port Gigabit Ethernet SFP Mini-PIM

Supported Platforms [SRX300, SRX320, SRX340, SRX550M](#)

Small form-factor pluggables (SFPs) are hot-pluggable modular interface transceivers for Gigabit and Fast Ethernet connections. Gigabit Ethernet SFP Mini-PIMs can be used in copper and optical environments to provide maximum flexibility when upgrading from an existing infrastructure to Metro Ethernet.

The 1-Port Gigabit Ethernet SFP Mini-PIM interfaces a single Gigabit Ethernet device or a network. It supports a variety of transceivers with data speeds of 10-Mbps/100-Mbps/1-Gbps with extended LAN or WAN connectivity.

Transceivers are hot-swappable.

This topic includes the following sections:

- [Supported Features on page 165](#)
- [Interface Names and Settings on page 166](#)
- [Available Link Speeds and Modes on page 166](#)
- [Link Settings on page 167](#)

Supported Features

The following features are supported on the 1-Port Gigabit Ethernet SFP Mini-PIM:

- 10-Mbps/100-Mbps/1-Gbps link speed
- Half-duplex/full-duplex support
- Autonegotiation

- Encapsulations
- Maximum transmission unit (MTU) size of 1514 bytes (default) and 9010 bytes (jumbo frames)
- Loopback
- Transceivers are hot-swappable

Interface Names and Settings

The following format is used to represent the 1-Port Gigabit Ethernet SFP Mini-PIM interface names:

type-fpc/pic/port

Where:

- **type**—Media type (**ge**)
- **fpc**—Number of the Flexible PIC Concentrator (FPC) card on which the physical interface is located
- **pic**—Number of the PIC on which the physical interface is located (0)
- **port**—Specific port on a PIC (0)

Examples: **ge-1/0/0** and **ge-2/0/0**

By default, the interfaces on the ports on the uplink module installed on the device are enabled. You can also specify the MTU size for the Gigabit Ethernet interface. Junos OS supports values from 256 through 9010. The default MTU size for Gigabit Ethernet interfaces is 1514.

Available Link Speeds and Modes

The 1-Port Gigabit Ethernet SFP Mini-PIM supports the following link speeds:

- **10m**—Sets the link speed to 10 Mbps.
- **100m**—Sets the link speed to 100 Mbps.
- **1g**—Sets the link speed to 1 Gbps.

The 1-Port Gigabit Ethernet SFP Mini-PIM supports the following link modes:

- **Full-duplex**—Allows bidirectional communication at a given point in time.
- **Half-duplex**—Allows single directional communication at a given point in time.

Link Settings

The 1-Port Gigabit Ethernet SFP Mini-PIM includes the following link settings:

- **auto-negotiation**—Enables autonegotiation of link mode and speed.



NOTE: By default, autonegotiation is enabled. To disable autonegotiation, use `set gigether-options no-autonegotiation`

We recommend enabling autonegotiation.

- **loopback**—Enables loopback.
- **no-auto-negotiation**—Disables autonegotiation of link mode and speed.
- **no-loopback**—Disables loopback.

By default a link speed of 1 Gbps in full-duplex mode is supported.



NOTE: On SRX340 High Memory devices, traffic might stop between the SRX340 device and the Cisco switch due to link mode mismatch. We recommend setting the same value to the autonegotiation parameters on both ends.



NOTE: On SRX300 devices, the link goes down when you upgrade FPGA on 1-Port Gigabit Ethernet SFP mini-PIM. As a workaround, run the `restart fpc` command and restart the FPC.

Related Documentation

- [Understanding Ethernet Interfaces on page 117](#)
- [Example: Configuring the 1-Port Gigabit Ethernet SFP Mini-PIM Interface on page 167](#)

Example: Configuring the 1-Port Gigabit Ethernet SFP Mini-PIM Interface

Supported Platforms [SRX300, SRX320, SRX340, SRX550M](#)

This example shows how to perform basic configuration for the 1-Port Gigabit Ethernet SFP Mini-PIM.

- [Requirements on page 168](#)
- [Overview on page 168](#)
- [Configuration on page 168](#)
- [Verification on page 171](#)

Requirements

Before you begin:

- Establish basic connectivity. See the *Getting Started Guide* for your device.
- Configure network interfaces as necessary. See “[Example: Creating an Ethernet Interface](#)” on page 122.

Overview

In this example, you configure the `ge-2/0/0` interface, set the operating speed to 100 Mbps, and define a logical interface that you can connect to the 1-Port Gigabit Ethernet SFP Mini-PIM. You also set the MTU value to 9010 and set the link option to no-loopback.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-2/0/0 link-mode full-duplex speed 100m
set interface ge-2/0/0 gigether-options no-loopback
```

Configuring Physical Properties

GUI Step-by-Step Procedure

To quickly configure the physical properties of a 1-Port Gigabit Ethernet SFP Mini-PIM using J-Web, use the following steps:

1. Select **Configure > Interfaces**.
2. Under Interface, select **ge-2/0/0** and then click **Edit**. A pop-up window appears.
3. In the Description box, type the description for the SFP Mini-PIM.
4. In the MTU box, type **9010**.
5. From the Speed list, select **100Mbps**.
6. From the Link-mode list, select **Full-duplex**.
7. Select the Enable Auto-negotiation checkbox.
8. Select the Enable Per Unit Scheduler checkbox.
9. Click **OK**.

Disabling the Interface

GUI Step-by-Step Procedure

To disable the 1-Port Gigabit Ethernet SFP Mini-PIM using J-Web, use the following steps:

1. Select **Configure > Interfaces**.
2. Under Interface, select **ge-2/0/0** and then click **Disable**.

Configuring Logical Properties

GUI Step-by-Step Procedure To quickly configure the logical properties of a 1-Port Gigabit Ethernet SFP Mini-PIM using J-Web, use the following steps:

1. Select **Configure > Interfaces**.
2. Under Interface, select **ge-2/0/0.0**, and then click **Add Logical Interface**. A pop-up window appears.
3. In the Unit box, type **0**.
4. In the Description box, type a description for the SFP Mini-PIM.
5. From the Zone list, select **untrust**.
6. To edit the family protocol type to the Mini-PIM interfaces, select the IPv4 tab, and then select **Enable address configuration**.
7. Click **Add**, and then type IPv4 address.
8. Click **OK**.

Editing Logical Properties

Step-by-Step Procedure To quickly configure the physical properties of a 1-Port Gigabit Ethernet SFP Mini-PIM using J-Web:

1. Under Interface, select the logical interface added to the 1-Port Gigabit Ethernet SFP Mini-PIM and then click **Edit**. A pop-up window appears.
2. Under Interface, select **ge-2/0/0.0**, and then click **Edit Logical Interface**. A pop-up window appears.
3. From the Zone list, select **trust**.
4. To enable DHCP client on the interface, select the IPv4 tab and then select **Enable DHCP**.
5. Click **OK**.



NOTE: You cannot add or edit Description and Unit for a logical interface.

Deleting the Logical Interface

GUI Step-by-Step Procedure To delete the logical interface of 1-Port Gigabit Ethernet SFP Mini-PIM using J-Web,

1. Select **Configure > Interfaces**.
2. Under Interface, select **ge-2/0/0.0**, and then click **Delete**.

Configuring a 1-Port Gigabit Ethernet SFP Mini-PIM

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a 1-Port Gigabit Ethernet SFP Mini-PIM:

1. Configure the interface.

```
[edit]
user@host# edit interfaces ge-2/0/0
```
2. Set the operating link-mode full-duplex speed of 100 Mbps for the SFP Mini-PIM.

```
[edit interfaces ge-2/0/0]
user@host# set link-mode full-duplex speed 100m
```
3. Assign the MTU value.

```
[edit interfaces ge-2/0/0]
user@host# set mtu 9010
```
4. Add the logical interface.

```
[edit interfaces ge-2/0/0]
user@host# set unit 0 family inet address 14.1.1.1/24
```
5. Set the link options.

```
[edit interfaces ge-2/0/0]
user@host# set gigheter-options no-loopback
```

Results From configuration mode, confirm your configuration by entering the **show interfaces ge-2/0/0** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces ge-2/0/0
mtu 9010;
speed 100m;
gigheter-options {
no-loopback;
}
unit 0 {
family inet {
14.1.1.1/24
}
}
```

If you are done configuring the device, enter **commit** from configuration mode.



NOTE: The 1-Port GE SFP Mini-PIM is installed in the second slot of the device chassis; therefore the output displayed is 1x GE High-Perf SFP mPIM and the Flexible PIC Concentrator (FPC) used here is fpc 2.

The 1-Port SFP Mini-PIM is installed in the fourth slot of the device chassis; therefore the output displayed is 1x GE SFP mPIM and Flexible PIC Concentrator (FPC) used here is fpc 4.

Verifying the FPC Status

Purpose Verify the FPC status.

Action From operational mode, enter the **show chassis fpc** command.

```
show@host> show chassis fpc
Slot State      Temp CPU Utilization (%)  Memory Utilization (%)
              (C) Total Interrupt  DRAM (MB) Heap      Buffer
0 Online       ----- CPU less FPC -----
1 Online       ----- CPU less FPC -----
2 Online       ----- CPU less FPC -----
3 Empty
4 Online       ----- CPU less FPC -----
```

The output should show the FPC status as online.

The 1-Port SFP Mini-PIM is installed in the fourth slot of the device chassis; the output shows the FPC status for slot 4 as online.

The 1-Port Gigabit Ethernet SFP Mini-PIM is installed in the second slot of the device chassis; the output shows the FPC status for slot 2 as online.

Verifying the Interface Settings

Purpose Verify that the interface is configured as expected.

Action From operational mode, enter the **show interface ge-2/0/0** command.

```
user@host# run show interfaces ge-2/0/0
Physical interface: ge-2/0/0, Enabled, Physical link is Up
Interface index: 156, SNMP ifIndex: 552
Link-level type: Ethernet, MTU: 9010, Link-mode: Full-duplex, Speed: 100mbps,
BPDU Error: None, MAC-REWRITE Error: None,
Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
Auto-negotiation: Enabled, Remote fault: Online
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags : None
CoS queues : 8 supported, 8 maximum usable queues
Current address: 00:22:83:99:ac:f2, Hardware address: 00:22:83:99:ac:f2
Last flapped : 2010-08-17 12:20:33 UTC (00:00:20 ago)
Input rate : 0 bps (0 pps)
```

```

Output rate      : 0 bps (0 pps)
Active alarms   : None
Active defects  : None

Logical interface ge-2/0/0.0 (Index 88) (SNMP ifIndex 557)
  Flags: SNMP-Traps Encapsulation: ENET2
  Input packets : 108
  Output packets: 1
  Security: Zone: Null
  Protocol inet, MTU: 8996
  Flags: Sendbcst-pkt-to-re
  Addresses, Flags: Is-Preferred Is-Primary
  Destination: 14.1.1/24, Local: 14.1.1.1, Broadcast: 14.1.1.255

```

Verify the following information in the command output:

- Physical interface—ge-2/0/0, Enabled, Physical link is Up
- MTU—9010; Link-mode—Full-duplex
- Speed—100 Mbps
- Loopback—Disabled

Related Documentation

- [Understanding Ethernet Interfaces on page 117](#)
- [Understanding the 1-Port Gigabit Ethernet SFP Mini-PIM on page 165](#)

Understanding the 2-Port 10-Gigabit Ethernet XPIM

Supported Platforms **SRX1500**

The 10-Gigabit Ethernet (also known as 10GBASE-T or IEEE 802.3an) is a telecommunication technology that offers data speeds up to 10 billion bits per second over unshielded or shielded twisted pair cables.

The 2-Port 10-Gigabit Ethernet Physical Interface Module (XPIM) is a 2 x 10GBASE-T / SFP+ XPIM line card. (SFP+ is a fiber optic transceiver module designed for 10-Gigabit Ethernet and 8.5 Gbps-fiber channel systems.) The 2-Port 10-Gigabit Ethernet XPIM provides a front-end interface connection that includes the following ports:

- 2 X copper ports. The copper ports support 10GBASE-T running with CAT6A or CAT7 Ethernet cable for up to 100 meters.
- 2 X fiber (SFP+) ports. The fiber ports support SFP+ multiple 10G modules.

The 2-Port 10-Gigabit Ethernet XPIM provides interconnects for LANs, WANs, and metropolitan area networks (MANs). The XPIM provides multiple service levels (1-Gigabit Ethernet to 10-Gigabit Ethernet in increments) and a single connection option for a wide range of customer needs and applications.



NOTE: By default, the 2-Port 10-Gigabit Ethernet XPIM ports comes up in fiber mode, while autonegotiation is not supported.

This topic includes the following sections:

- [Supported Features on page 174](#)
- [Interface Names and Settings on page 174](#)
- [Copper and Fiber Operating Modes on page 175](#)
- [Link Speeds on page 175](#)
- [Link Settings on page 175](#)

Supported Features

The following features are supported on the 2-Port 10-Gigabit Ethernet XPIM:

- Multiple SFP+ 10G modules and the following SFP modules:
 - SFPP-10GE-SR
 - SFPP-10GE-LR
 - SFPP-10GE-ER
 - SFPP-10GE-LRM
- Copper TWIN-AX 1M and Copper TWIN-AX 3M
- Online Insertion and Removal (OIR) functionality
- Link speeds of up to 10-Gbps
- Full-duplex and half-duplex modes
- Flow control
- Autonegotiation and autosensing
- Quality of service (QoS)

Interface Names and Settings

The following format is used to represent the 2-Port 10-Gigabit Ethernet XPIM interface names:

type-fpc/pic/port

Where:

- type — Media type (xe)
- fpc — Number of the Flexible PIC Concentrator (FPC) card on which the physical interface is located
- pic — Number of the PIC on which the physical interface is located (0)
- port — Specific port on a PIC (0 or 1)

By default, the interfaces (for example, **xe-6/0/0** or **xe-2/0/0**) on the ports on the uplink module installed on the device are enabled. You can also specify the maximum

transmission unit (MTU) size for the Gigabit Ethernet interface. Junos OS supports values from 256 through 9192. The default MTU for Gigabit Ethernet interfaces is 1514.

Copper and Fiber Operating Modes

On the 2-Port 10-Gigabit Ethernet XPIM, one copper port and one fiber port is grouped together as port 0, and another copper port and fiber port are grouped as port 1. Only two ports can be active at the same time (one port from port 0 and another port from port 1).

The 2-Port 10-Gigabit Ethernet XPIM can be configured to operate in two copper mode, two fiber mode, or mixed mode (one copper and one fiber). In mixed mode, the two ports should be from different port groups (one port from port 1 and the other from port 2).

Link Speeds

The 2-Port 10-Gigabit Ethernet XPIM ports support the following link speeds for copper and fiber:

- Copper—10/100/1000 Mbps or 10Gbps (full duplex). Half-duplex is only for 10/100 Mbps.
- Fiber—1000 Mbps or 10 Gbps (full duplex). Half-duplex mode is not supported.

To set the link speeds, use the following options:

- **10m**—Sets the link speed to 10 Mbps.
- **10g**—Sets the link speed to 10 Gbps.
- **100m**—Sets the link speed to 100 Mbps.
- **1g**—Sets the link speed to 1 Gbps.

Link Settings

The 2-Port 10-Gigabit Ethernet XPIM includes the following link settings:

- **802.3ad**—Specifies an aggregated Ethernet bundle.
- **auto-negotiation**—Enables autonegotiation of flow control, link mode, and speed.
- **loopback**—Enables loopback.
- **no-auto-negotiation**—Disables autonegotiation of flow control, link mode, and speed.
- **no-loopback**—Disables loopback.

By default, flow control is enabled on all ports, a link speed of 10 Gbps in full duplex is supported, autonegotiation is disabled on the fiber ports, and autonegotiation is enabled on copper ports.



NOTE: Autonegotiation is not supported when the 2-Port 10-Gigabit Ethernet XPIM is operating in fiber mode at a link speed of 10 Gbps.

- Related Documentation**
- [Understanding Ethernet Interfaces on page 117](#)
 - [Example: Configuring the 2-Port 10-Gigabit Ethernet XPIM Interface on page 176](#)

Example: Configuring the 2-Port 10-Gigabit Ethernet XPIM Interface

Supported Platforms [SRX5400, SRX5600, SRX5800](#)

This example shows how to perform basic configuration for the 1-Port Gigabit Ethernet SFP Mini-PIM.

- [Requirements on page 176](#)
- [Overview on page 176](#)
- [Configuration on page 176](#)
- [Verification on page 178](#)

Requirements

Before you begin:

- Establish basic connectivity. See the *Getting Started Guide* for your device.
- Configure network interfaces as necessary. See “[Example: Creating an Ethernet Interface](#)” on page 122.

Overview

In this example, you configure the xe-6/0/0 interface, set the operating mode to copper mode, set the operating speed to 10 Gbps, and define a logical interface that you can connect to the 2-Port 10-Gigabit Ethernet XPIM. Additionally, you set the MTU value to 1514, set the link option to no loopback, and enable the interface.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces xe-6/0/0 media-type copper speed 10g unit 0 family inet mtu 1514
set interface xe-6/0/0 gigether-options no-loopback
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a 2-Port 10-Gigabit Ethernet XPIM:

1. Configure the interface.

```
[edit]
user@host# edit interfaces xe-6/0/0
```

2. Configure the operating mode.


```
[edit interfaces xe-6/0/0]
user@host# set media-type copper
```
3. Set the operating speed for the XPIM.


```
[edit interfaces xe-6/0/0]
user@host# set speed 10g
```
4. Add the logical interface.


```
[edit interfaces xe-6/0/0]
user@host# set unit 0 family inet
```
5. Assign the physical interface MTU value.


```
[edit interfaces xe-6/0/0]
user@host# set interface xe-6/0/0 mtu 1514
```
6. Assign the logical interface MTU value.


```
[edit interfaces xe-6/0/0]
user@host# set unit 0 family inet mtu 1500
```
7. Set the link options.


```
[edit interfaces xe-6/0/0]
user@host# set gigether-options no-loopback
```
8. Disable the interface.


```
[edit interfaces xe-6/0/0]
user@host# set disable
```
9. Enable the interface.


```
[edit interfaces xe-6/0/0]
user@host# delete disable
```

Results From configuration mode, confirm your configuration by entering the **show interfaces xe-6/0/0** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces xe-6/0/0
speed 10g;
media-type copper;
gigether-options {
no-loopback;
}
unit 0 {
family inet {
mtu 1514;
}
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying That the Correct Hardware Is Installed on page 178](#)
- [Verifying the FPC Status on page 178](#)
- [Verifying the Interface Settings on page 179](#)

Verifying That the Correct Hardware Is Installed

Purpose Verify that the 2-Port 10-Gigabit Ethernet XPIM is installed on the device.

Action From operational mode, enter the **show chassis hardware** command.

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis		AJ0309AC0047	SRX650	
Midplane	REV 04	710-023875	TV3993	
System IO	REV 04	710-023209	TV4035	SRXSME System IO
Routing Engine	REV 01	710-023224	DT5109	RE-SRXSME-SRE6
FPC 0				FPC
PIC 0				4x GE Base PIC
FPC 2				FPC
PIC 0				2x 10G gPIM
FPC 6				FPC
PIC 0				2x 10G gPIM
Power Supply 0	REV 01	740-024283	TA00049WSSSS	PS 645W AC

Verify that the output contains the following values:

- FPC 2, PIC 0—2x 10G gPIM
- FPC 6, PIC 0—2x 10G gPIM

Verifying the FPC Status

Purpose Verify the FPC status.

Action From operational mode, enter the **show chassis fpc** command.

Slot	Temp	CPU Utilization (%)	Memory Utilization (%)
State	(C)	Total	Interrupt
		DRAM (MB)	Heap Buffer
0	Online	-----	CPU less FPC -----
1	Empty		
2	Online	-----	CPU less FPC -----
3	Empty		
4	Empty		
5	Empty		
6	Online	-----	CPU less FPC -----
7	Empty		
8	Empty		

The output should display FPC status as online.

Verifying the Interface Settings

Purpose Verify that the interface is configured as expected.

Action From operational mode, enter the **show interface xe-6/0/0** command.

```
Physical interface: xe-6/0/0, Enabled, Physical link is Up
Interface index: 144, SNMP ifIndex: 501
Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 10Gbps,
BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled
Device flags : Present Running
6 Copyright © 2010, Juniper Networks, Inc.
Interface flags: SNMP-Traps Internal: 0x0
Link flags : None
CoS queues : 8 supported, 8 maximum usable queues
Current address: 00:1f:12:e0:80:a8, Hardware address: 00:1f:12:e0:80:a8
Last flapped : 1970-01-01 00:34:22 PST (07:26:29 ago)
Input rate : 0 bps (0 pps)
Output rate : 0 bps (0 pps)
Active alarms : None
Active defects : None
```

```
Logical interface xe-6/0/0.0 (Index 72) (SNMP ifIndex 503)
Flags: SNMP-Traps Encapsulation: ENET2
Input packets : 25
Output packets: 25
Security: Zone: HOST
Allowed host-inbound traffic : any-service bfd bgp dvmrp igmp ldp msdp nhrp
ospf pgm pim rip router-discovery rsvp sap vrrp
Protocol inet, MTU: 1500
Flags: Sendbcst-pkt-to-re
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.10.10/24, Local: 10.10.10.10, Broadcast: 10.10.10.255
```

Verify the following information in the command output:

- Physical interface—xe-6/0/0, Enabled, Physical link is Up
- MTU—1514
- Link mode—Full duplex
- Speed—10 Gbps
- Loopback—Disabled
- Flow control—Enabled

Related Documentation

- [Understanding the 2-Port 10-Gigabit Ethernet XPIM on page 173](#)
- [Understanding Ethernet Interfaces on page 117](#)

Configuring Ethernet OAM Link Fault Management

- [Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways on page 181](#)
- [Example: Configuring Ethernet OAM Link Fault Management on page 183](#)

Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, SRX550M](#)

The Ethernet interfaces on SRX Series devices support the IEEE 802.3ah standard for Operation, Administration, and Maintenance (OAM). The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters. The IEEE 802.3ah standard meets the requirement for OAM capabilities as Ethernet moves from being solely an enterprise technology to a WAN and access technology, and the standard remains backward-compatible with existing Ethernet technology.

The following OAM LFM features are supported:

- **Discovery and link monitoring**—The discovery process is triggered automatically when OAM is enabled on the interface. The discovery process permits Ethernet interfaces to discover and monitor the peer on the link if it also supports the IEEE 802.3ah standard. In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality. In passive mode, the peer initiates the discovery process. After the discovery process has been initiated, both sides participate in discovery. The device performs link monitoring by sending periodic OAM protocol data units (PDUs) to advertise OAM mode, configuration, and capabilities.

You can specify the number of OAM PDUs that an interface can miss before the link between peers is considered down.

- **Remote fault detection**—Remote fault detection uses flags and events. Flags convey Link Fault (a loss of signal), Dying Gasp (an unrecoverable condition such as a power failure), and Critical Event (an unspecified vendor-specific critical event). You can specify the periodic OAM PDU sending interval for fault detection. SRX Series devices use the Event Notification OAM PDU to notify the remote OAM device when a problem

is detected. You can specify the action to be taken by the system when the configured link-fault event occurs.

- Remote loopback—Remote loopback mode ensures link quality between the device and a remote peer during installation or troubleshooting. In this mode, when the interface receives a frame that is not an OAM PDU or a pause frame, it sends it back on the same interface on which it was received. The link appears to be in the active state. You can use the returned loopback acknowledgement to test delay, jitter, and throughput.

Junos OS can place a remote data terminal equipment (DTE) into loopback mode (if remote loopback mode is supported by the remote DTE). When you place a remote DTE into loopback mode, the interface receives the remote loopback request and puts the interface into remote loopback mode. When the interface is in remote loopback mode, all frames except OAM PDUs are looped back without any changes made to the frames. OAM PDUs continue to be sent and processed.

Table 18 on page 182 lists the interfaces modes supported.

Table 18: Supported Interface Modes

Interfaces	Mode
Physical interface (fe/ge)	Family <ul style="list-style-type: none"> • ccc • ethernet-switching • inet6 • inet • iso • mpls • tcc
	IFD encapsulations <ul style="list-style-type: none"> • ethernet-ccc • extended-vlan-ccc (IFD vlan-tagging mode) • ethernet-tcc • extended-vlan-tcc
Aggregated Ethernet interface (Static or LACP lag)	Family <ul style="list-style-type: none"> • ethernet-switching • inet • mpls • iso • inet6
	IFD encapsulations <ul style="list-style-type: none"> • ethernet-ccc • extended-vlan-ccc (IFD vlan-tagging mode) • vlan-ccc

- Related Documentation**
- [Example: Configuring Ethernet OAM Link Fault Management on page 183](#)

Example: Configuring Ethernet OAM Link Fault Management

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, SRX550M](#)

The Ethernet interfaces on the SRX Series devices support the IEEE 802.3ah standard for Operation, Administration, and Maintenance (OAM). The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters.

This example describes how to enable and configure OAM LFM on a Gigabit Ethernet or Fast Ethernet interface:

- [Requirements on page 183](#)
- [Overview on page 183](#)
- [Configuration on page 184](#)
- [Verification on page 186](#)

Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.1 R2 or later for SRX Series Services Gateways
- Any two models of SRX Series devices connected directly

Before you begin:

- Establish basic connectivity. See the Getting Started Guide for your device.
- Configure network interfaces as necessary. See [“Example: Creating an Ethernet Interface” on page 122](#).
- Ensure that you configure the interfaces as per the interface modules listed in [“Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways” on page 181](#)

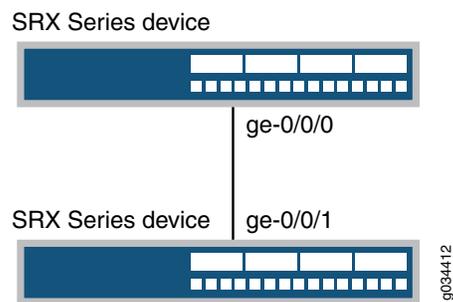
Overview

The Ethernet interfaces on the SRX Series devices support the IEEE 802.3ah standard for Operation, Administration, and Maintenance (OAM). The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters.

This example uses two SRX Series devices connected directly. Before you begin configuring Ethernet OAM LFM on these two devices, connect the two devices directly through supported interfaces. See [“Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways” on page 181](#).

[Figure 11 on page 184](#) shows the topology used in this example.

Figure 11: Ethernet LFM with SRX Series Devices



NOTE: For more information about configuring Ethernet OAM Link Fault Management, see [Junos® OS Ethernet Interfaces](#).

Configuration

To configure Ethernet OAM LFM, perform these tasks:

- [Configuring Ethernet OAM Link Fault Management on Device 1 on page 184](#)
- [Configuring Ethernet OAM Link Fault Management on Device 2 on page 185](#)

Configuring Ethernet OAM Link Fault Management on Device 1

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set protocols oam ethernet link-fault-management interface ge-0/0/0
set protocols oam ethernet link-fault-management interface ge-0/0/0 link-discovery
  active
set protocols oam ethernet link-fault-management interface ge-0/0/0 pdu-interval 800
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the Junos OS CLI User Guide.

To configure Ethernet OAM LFM on device 1:

1. Enable IEEE 802.3ah OAM support.


```
[edit protocols oam ethernet link-fault-management]
user@device1# set interface ge-0/0/0
```
2. Specify that the interface initiates the discovery process.


```
[edit protocols oam ethernet link-fault-management]
user@device1# set interface ge-0/0/0 link-discovery active
```
3. Set the periodic OAM PDU-sending interval (in milliseconds) for fault detection.


```
[edit protocols oam ethernet link-fault-management]
user@device1# set interface pdu-interval 800
```

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

```
[edit]
user@device1# show protocols
protocols {
  oam {
    ethernet {
      link-fault-management {
        interface ge-0/0/0 {
          pdu-interval 800;
          link-discovery active;
        }
      }
    }
  }
}
```

Configuring Ethernet OAM Link Fault Management on Device 2

CLI Quick Configuration To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set protocols oam ethernet link-fault-management interface ge-0/0/1
set protocols oam ethernet link-fault-management interface ge-0/0/1 pdu-interval 800
set protocols oam ethernet link-fault-management interface ge-0/0/1 negotiation-options
allow-remote-loopback
```

Step-by-Step Procedure To configure Ethernet OAM LFM on device 2:

1. Enable OAM on the peer interface.


```
[edit protocols oam ethernet link-fault-management]
user@device2# set interface ge-0/0/1
```
2. Set the periodic OAM PDU-sending interval (in milliseconds) for fault detection.


```
[edit protocols oam ethernet link-fault-management]
user@device2# set interface ge-0/0/1 pdu-interval 800
```
3. Enable remote loopback support for the local interface.


```
[edit protocols oam ethernet link-fault-management]
user@device2# set interface ge-0/0/1 negotiation-options allow-remote-loopback
```

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

```
[edit]
user@device2# show protocols
protocols {
  oam {
```

```

ethernet {
  link-fault-management {
    interface ge-0/0/1 {
      negotiation-options {
        allow-remote-loopback;
      }
    }
  }
}

```

Verification

Verify the OAM LFM Configuration

- Purpose** Verify that OAM LFM is configured properly.
- Action** From operational mode, enter the `show oam ethernet link-fault-management` command.
- ```

user@device1>show oam ethernet link-fault-management

Interface: ge-0/0/0.0
Status: Running, Discovery state: Send Any
Peer address: 00:19:e2:50:3b:e1
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
Remote entity information:
Remote MUX action: forwarding, Remote parser action: forwarding
Discovery mode: active, Unidirectional mode: unsupported
Remote loopback mode: supported, Link events: supported
Variable requests: unsupported

```
- Meaning** The output displays the MAC address and the discovery state is **Send Any** if OAM LFM has been configured properly.
- Related Documentation**
- [Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways on page 181](#)

# Configuring Power over Ethernet

- [Understanding Power over Ethernet on page 187](#)
- [Example: Configuring PoE on All Interfaces on page 189](#)
- [Example: Configuring PoE on an Individual Interface on page 191](#)
- [Example: Disabling a PoE Interface on page 194](#)

## Understanding Power over Ethernet

---

### Supported Platforms [SRX1500, SRX320, SRX340](#)

Power over Ethernet (PoE) is the implementation of the IEEE 802.3 AF and IEEE 802.3 AT standards that allow both data and electrical power to pass over a copper Ethernet LAN cable.

The SRX Series devices support PoE on Ethernet ports. PoE ports transfer electrical power and data to remote devices over standard twisted-pair cable in an Ethernet network. PoE ports allow you to plug in devices that require both network connectivity and electrical power, such as VoIP and IP phones and wireless LAN access points.

You can configure the SRX Series device to act as power sourcing equipment (PSE), supplying power to powered devices that are connected on designated ports.

This topic contains the following sections:

- [SRX Series Services Gateway PoE Specifications on page 187](#)
- [PoE Classes and Power Ratings on page 188](#)
- [PoE Options on page 189](#)

### SRX Series Services Gateway PoE Specifications

[Table 19 on page 187](#) lists the PoE specifications for the SRX320 and SRX340 devices

**Table 19: PoE Specifications for the SRX320 and SRX340 Devices**

| Specifications      | For SRX320 Device                                                                                   | For SRX340 Device                                                                                                                   |
|---------------------|-----------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------|
| Supported standards | <ul style="list-style-type: none"> <li>• IEEE 802.3 AF</li> <li>• Legacy (pre-standards)</li> </ul> | <ul style="list-style-type: none"> <li>• IEEE 802.3 AF</li> <li>• IEEE 802.3 AT (PoE+)</li> <li>• Legacy (pre-standards)</li> </ul> |

**Table 19: PoE Specifications for the SRX320 and SRX340 Devices (continued)**

| Specifications                    | For SRX320 Device                                                                                                                                                                                                     | For SRX340 Device                                                                                                                                                                                                     |
|-----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Supported ports                   | Supported on two Gigabit Ethernet ports and two Fast Ethernet ports (ge-0/0/0, ge-0/0/1, fe-0/0/2, and fe-0/0/3).                                                                                                     | Supported on all 16 Gigabit Ethernet ports (ge-0/0/0 to ge-0/0/15).                                                                                                                                                   |
| Total PoE power sourcing capacity | 50 W                                                                                                                                                                                                                  | 150 W                                                                                                                                                                                                                 |
| Default per port power limit      | 15.4 W                                                                                                                                                                                                                | 15.4 W                                                                                                                                                                                                                |
| Maximum per port power limit      | 30 W                                                                                                                                                                                                                  | 30 W                                                                                                                                                                                                                  |
| Power management modes            | <ul style="list-style-type: none"> <li>• Static: Power allocated for each interface can be configured.</li> <li>• Class: Power allocated for interfaces is based on the class of powered device connected.</li> </ul> | <ul style="list-style-type: none"> <li>• Static: Power allocated for each interface can be configured.</li> <li>• Class: Power allocated for interfaces is based on the class of powered device connected.</li> </ul> |

### PoE Classes and Power Ratings

A powered device is classified based on the maximum power that it draws across all input voltages and operational modes. When class-based power management mode is configured on the SRX Series devices, power is allocated taking into account the maximum power ratings defined for the different classes of devices.

Table 20 on page 188 lists the classes and their power ratings as specified by the IEEE standards.

**Table 20: SRX Series Devices PoE Specifications**

| Class | Usage    | Minimum Power Levels Output from PoE Port                                                   |
|-------|----------|---------------------------------------------------------------------------------------------|
| 0     | Default  | 15.4 W                                                                                      |
| 1     | Optional | 4.0 W                                                                                       |
| 2     | Optional | 7.0 W                                                                                       |
| 3     | Optional | 15.4 W                                                                                      |
| 4     | Reserved | Class 4 power devices are eligible to receive power up to 30 W according to IEEE standards. |

## PoE Options

When configuring PoE, you must enable the PoE interface in order for the port to provide power to a connected, powered device. In addition, you can configure the following PoE features:

- Port priority—Sets port priority. When it is not possible to maintain power to all connected ports, lower priority ports are powered off before higher priority ports. When a new device is connected on a higher-priority port, a lower priority port will be powered off automatically if available power is insufficient to power on the higher priority port. (For the ports with the same priority configuration, ports on the left are given higher priority than the ports on the right.)
- Maximum available wattage power available to a port—Sets the maximum amount of power that can be supplied to the port. The default wattage per port is 15.4 watts.
- PoE power consumption logging—Allows logging of per-port PoE power consumption. The telemetries section must be explicitly specified to enable logging. If left unspecified, telemetries is disabled by default. The default telemetry duration is 1 hour. The default telemetry interval is 5 minutes.
- PoE power management mode—Has two modes:
  - Class—When a powered device is connected to a PoE port, the power allocated to it is equal to the maximum power for the class as defined by the IEEE standards.
  - Static—When a powered device is connected to a PoE port, the power allocated to it is equal to the maximum power configured for the port.
- Reserve power—Reserves the specified amount of power for the gateway in case of a spike in PoE consumption. The default is 0.

### Related Documentation

- [Understanding Ethernet Interfaces on page 117](#)
- [Example: Configuring PoE on All Interfaces on page 189](#)
- [Example: Configuring PoE on an Individual Interface on page 191](#)
- [Example: Disabling a PoE Interface on page 194](#)

## Example: Configuring PoE on All Interfaces

**Supported Platforms** SRX1500, SRX320, SRX340

This example shows how to configure PoE on all interfaces.

- [Requirements on page 190](#)
- [Overview on page 190](#)
- [Configuration on page 190](#)
- [Verification on page 191](#)

## Requirements

Before you begin, configure Ethernet interfaces. See [“Example: Creating an Ethernet Interface” on page 122](#).

## Overview

This example shows how to configure PoE on all interfaces on a device. In this example, you set the power port priority to low and the maximum power available to a port to 15.4 watts. Then you enable the PoE power consumption logging with the default telemetries settings, and you set the PoE management mode to static. Finally, you set the reserved power consumption to 15 watts in case of a spike in PoE consumption.

## Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set poe interface all priority low maximum-power 15.4 telemetries
set poe management static guard-band 15
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see the [Junos OS CLI User Guide](#).

To configure PoE on all interfaces:

1. Enable PoE.  

```
[edit]
user@host# edit poe interface all
```
2. Set the power port priority.  

```
[edit poe interface all]
user@host# set priority low
```
3. Set the maximum PoE wattage available for a port.  

```
[edit poe interface all]
user@host# set maximum-power 15.4
```
4. Enable logging of PoE power consumption.  

```
[edit poe interface all]
user@host# set telemetries
```
5. Set the PoE management mode.  

```
[edit]
user@host# set poe management static
```
6. Reserve power wattage in case of a spike in PoE consumption.  

```
[edit]
user@host# set poe guard-band 15
```

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

```
[edit]
user@host# show poe interface all
priority low;
maximum-power 15.4;
telemetries;
```

If you are done configuring the device, enter **commit** from configuration mode.

## Verification

Confirm that the configuration is working properly.

### Verifying the Status of PoE Interfaces

**Purpose** Verify that the PoE interfaces on the device are enabled and set to the desired priority settings. (The device used here is the SRX340 Services Gateway.)

**Action** From operational mode, enter the **show poe interface all** command.

```
user@host> show poe interface all

Interface Admin status Oper status Max power Priority Power consumption Class
ge-0/0/0 Enabled Searching 15.4W Low 0.0W 0
ge-0/0/1 Enabled Powered-up 15.4W High 6.6W 0
ge-0/0/2 Disabled Disabled 15.4W Low 0.0W 0
ge-0/0/3 Disabled Disabled 15.4W Low 0.0W 0
```

The **show poe interface all** command lists PoE interfaces configured on the SRX 240 device, including information on status, priority, power consumption, and class. This output shows that the device has four PoE interfaces of which two are enabled with default values. One port has a device connected that is drawing power within expected limits.

- Related Documentation**
- [Understanding Power over Ethernet on page 187](#)
  - [Example: Configuring PoE on an Individual Interface on page 191](#)
  - [Example: Disabling a PoE Interface on page 194](#)

## Example: Configuring PoE on an Individual Interface

**Supported Platforms** [SRX1500](#), [SRX210](#), [SRX220](#), [SRX240](#)

This example shows how to configure PoE on an individual interface.

- [Requirements on page 192](#)
- [Overview on page 192](#)

- [Configuration on page 192](#)
- [Verification on page 193](#)

## Requirements

Before you begin:

- Configure Ethernet interfaces. See [“Example: Creating an Ethernet Interface” on page 122](#).
- Configure PoE on all interfaces. See [“Example: Configuring PoE on All Interfaces” on page 189](#).

## Overview

This example shows how to configure PoE on the ge-0/0/0 interface. In this example, you set the power port priority to high and the maximum power available to a port to 15.4 watts. Then you enable the PoE power consumption logging with the default telemetries settings, and you set the PoE management mode to static. Finally, you set the reserved power to 15 watts in case of a spike in PoE consumption.

## Configuration

### CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set poe interface ge-0/0/0 priority high maximum-power 15.4 telemetries
set poe management static guard-band 15
```

### Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see the [Junos OS CLI User Guide](#).

To configure PoE:

1. Enable PoE.  

```
[edit]
user@host# edit poe interface ge-0/0/0
```
2. Set the power port priority.  

```
[edit poe interface ge-0/0/0]
user@host# set priority high
```
3. Set the maximum PoE wattage available for a port.  

```
[edit poe interface ge-0/0/0]
user@host# set maximum power 15.4
```
4. Enable logging of PoE power consumption.  

```
[edit poe interface ge-0/0/0]
user@host# set telemetries
```
5. Set the PoE management mode.  

```
[edit]
```

```
user@host# set poe management static
```

- Reserve power wattage in case of a spike in PoE consumption.

```
[edit]
user@host# set poe guard-band 15
```

**Results** From configuration mode, confirm your configuration by entering the **show poe interface ge-0/0/0** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show poe interface ge-0/0/0
 priority high;
 maximum-power 15.4;
 telemetries;
```

If you are done configuring the device, enter **commit** from configuration mode.

## Verification

Confirm that the configuration is working properly.

- [Verifying the Status of PoE Interfaces on page 193](#)
- [Verifying the Telemetry Data \(History\) for the Specified Interface on page 193](#)
- [Verifying PoE Global Parameters on page 194](#)

### Verifying the Status of PoE Interfaces

**Purpose** Verify that the PoE interfaces on the device are enabled and set to the desired priority settings. (The device used in this example is the SRX340 Services Gateway.)

**Action** From operational mode, enter the **show poe interface ge-0/0/1** command.

```
user@host> show poe interface ge-0/0/1
PoE interface status:
PoE interface : ge-0/0/1
Administrative status : Enabled
Operational status : Powered-up
Power limit on the interface : 15.4 W
Priority : High
Power consumed : 6.6 W
Class of power device : 0
```

The **show poe interface ge-0/0/1** command lists PoE interfaces configured on the SRX340 device, with their status, priority, power consumption, and class.

### Verifying the Telemetry Data (History) for the Specified Interface

**Purpose** Verify the PoE interface's power consumption over a specified period.

**Action** From operational mode, enter the **show poe telemetries interface** command.

For all records:

```
user@host> show poe telemetries interface ge-0/0/1 all
SI No Timestamp Power Voltage
1 Fri Jan 04 11:41:15 2009 5.1 W 47.3 V
2 Fri Jan 04 11:40:15 2009 5.1 W 47.3 V
3 Fri Jan 04 11:39:15 2009 5.1 W 47.3 V
4 Fri Jan 04 11:38:15 2009 0.0 W 0.0 V
5 Fri Jan 04 11:37:15 2009 0.0 W 0.0 V
6 Fri Jan 04 11:36:15 2009 6.6 W 47.2 V
7 Fri Jan 04 11:35:15 2009 6.6 W 47.2 V
```

For a specific number of records:

```
user@host> show poe telemetries interface ge-0/0/15
SI No Timestamp Power Voltage
1 Fri Jan 04 11:31:15 2009 6.6 W 47.2 V
2 Fri Jan 04 11:30:15 2009 6.6 W 47.2 V
3 Fri Jan 04 11:29:15 2009 6.6 W 47.2 V
4 Fri Jan 04 11:28:15 2009 6.6 W 47.2 V
5 Fri Jan 04 11:27:15 2009 6.6 W 47.2 V
```

The telemetry status displays the power consumption history for the specified interface, provided telemetry has been configured for that interface.

### Verifying PoE Global Parameters

---

**Purpose** Verify global parameters such as guard band, power limit, and power consumption.

**Action** From operational mode, enter the **show poe controller** command.

```
user@host> show poe controller
Controller Maximum Power Guard band Management
index power consumption 0 W Static
0 150.0 W 0.0 W 0 W Static
```

The **show poe controller** command lists the global parameters configured on the SRX Series device such as controller index, maximum power, power consumption, guard band, and management mode along with their status.

- Related Documentation**
- [Understanding Power over Ethernet on page 187](#)
  - [Example: Configuring PoE on All Interfaces on page 189](#)
  - [Example: Disabling a PoE Interface on page 194](#)

### Example: Disabling a PoE Interface

---

**Supported Platforms** SRX1500, SRX320, SRX340

This example shows how to disable PoE on all interfaces or on a specific interface.

- [Requirements on page 195](#)
- [Overview on page 195](#)

- [Configuration on page 195](#)
- [Verification on page 195](#)

## Requirements

Before you begin:

- Configure PoE on all interfaces. See “[Example: Configuring PoE on All Interfaces](#)” on [page 189](#).
- Configure PoE on an individual interface. See “[Example: Configuring PoE on an Individual Interface](#)” on [page 191](#).

## Overview

In this example, you disable PoE on all interfaces and on a specific interface, which in this case is ge-0/0/0.

## Configuration

### Step-by-Step Procedure

To disable PoE on interfaces:

1. Disable PoE on all interfaces.  
[edit]  
user@host# **set poe interface all disable**
2. Disable PoE on a specific interface.  
[edit]  
user@host# **set poe interface ge-0/0/0 disable**
3. If you are done configuring the device, commit the configuration.  
[edit]  
user@host# **commit**

## Verification

To verify the configuration is working properly, enter the **show poe interface** command.

### Related Documentation

- [Understanding Power over Ethernet on page 187](#)



## PART 5

# Configuring Interface Encapsulation

- [Interface Encapsulation Overview on page 199](#)
- [Configuring Point-to-Point Protocol over Ethernet on page 207](#)
- [Configuring PPPoE-Based Radio-to-Router Protocol on page 229](#)
- [Configuring R2CP Radio-to-Router Protocol on page 237](#)



# Interface Encapsulation Overview

- [Understanding Physical Encapsulation on an Interface on page 199](#)
- [Understanding Frame Relay Encapsulation on an Interface on page 200](#)
- [Understanding Point-to-Point Protocol on page 202](#)
- [Understanding High-Level Data Link Control on page 204](#)

## Understanding Physical Encapsulation on an Interface

---

**Supported Platforms** [SRX1500, SRX300, SRX320, SRX340](#)

Encapsulation is the process by which a lower level protocol accepts a message from a higher level protocol and places it in the data portion of the lower level frame. As a result, datagrams transmitted through a physical network have a sequence of headers: the first header for the physical network (or Data Link Layer) protocol, the second header for the Network Layer protocol (IP, for example), the third header for the Transport Layer protocol, and so on.

The following encapsulation protocols are supported on physical interfaces:

- [Frame Relay Encapsulation](#). See “[Understanding Frame Relay Encapsulation on an Interface](#)” on [page 200](#).
- [Point-to-Point Protocol](#). See “[Understanding Point-to-Point Protocol](#)” on [page 202](#).
- [Point-to-Point Protocol over Ethernet](#). See “[Understanding Point-to-Point Protocol over Ethernet](#)” on [page 207](#).
- [High-Level Data Link Control](#). See “[Understanding High-Level Data Link Control](#)” on [page 204](#).

**Related Documentation**

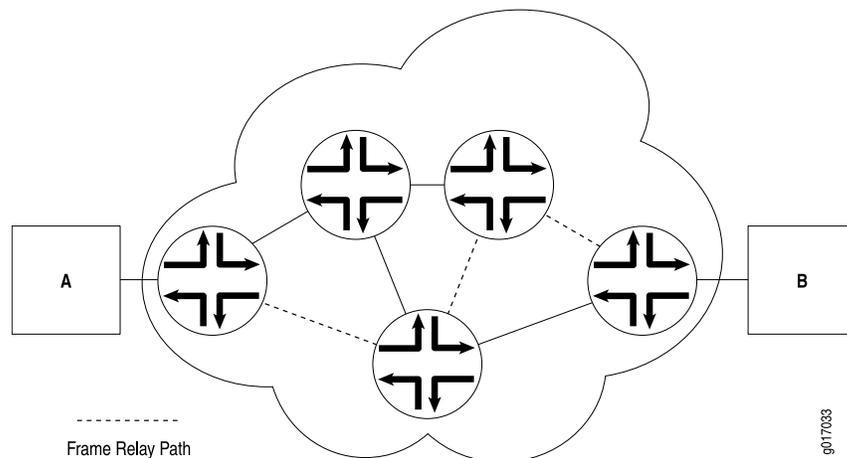
- [Understanding Interfaces on page 3](#)
- [Understanding Frame Relay Encapsulation on an Interface on page 200](#)
- [Understanding Point-to-Point Protocol on page 202](#)
- [Understanding High-Level Data Link Control on page 204](#)

## Understanding Frame Relay Encapsulation on an Interface

**Supported Platforms** SRX1500, SRX320, SRX340

The Frame Relay packet-switching protocol operates at the Physical Layer and Data Link Layer in a network to optimize packet transmissions by creating virtual circuits between hosts. [Figure 12 on page 200](#) shows a typical Frame Relay network.

**Figure 12: Frame Relay Network**



[Figure 12 on page 200](#) shows multiple paths from Host A to Host B. In a typical routed network, traffic is sent from device to device with each device making routing decisions based on its own routing table. In a packet-switched network, the paths are predefined. Devices switch a packet through the network according to predetermined next-hops established when the virtual circuit is set up.

This topic contains the following sections:

- [Virtual Circuits on page 200](#)
- [Switched and Permanent Virtual Circuits on page 201](#)
- [Data-Link Connection Identifiers on page 201](#)
- [Congestion Control and Discard Eligibility on page 201](#)

### Virtual Circuits

A virtual circuit is a bidirectional path between two hosts in a network. Frame Relay virtual circuits are logical connections between two hosts that are established either by a call setup mechanism or by an explicit configuration.

A virtual circuit created through a call setup mechanism is known as a switched virtual circuit (SVC). A virtual circuit created through an explicit configuration is called a permanent virtual circuit (PVC).

## Switched and Permanent Virtual Circuits

Before data can be transmitted across an SVC, a signaling protocol like ISDN must set up a call by the exchange of setup messages across the network. When a connection is established, data is transmitted across the SVC. After data transmission, the circuit is torn down and the connection is lost. For additional traffic to pass between the same two hosts, a subsequent SVC must be established, maintained, and terminated.

Because PVCs are explicitly configured, they do not require the setup and teardown of SVCs. Data can be switched across the PVC whenever a host is ready to transmit. SVCs are useful in networks where data transmission is sporadic and a permanent circuit is not needed.

## Data-Link Connection Identifiers

An established virtual circuit is identified by a data-link connection identifier (DLCI). The DLCI is a value from 16 through 1022. (Values 1 through 15 are reserved.) The DLCI uniquely identifies a virtual circuit locally so that devices can switch packets to the appropriate next-hop address in the circuit. Multiple paths that pass through the same transit devices have different DLCIs and associated next-hop addresses.

## Congestion Control and Discard Eligibility

Frame Relay uses the following types of congestion notification to control traffic within a Frame Relay network. Both are controlled by a single bit in the Frame Relay header.

- Forward explicit congestion notification (FECN)
- Backward explicit congestion notification (BECN)

Traffic congestion is typically defined in the buffer queues on a device. When the queues reach a predefined level of saturation, traffic is determined to be congested. When traffic congestion occurs in a virtual circuit, the device experiencing congestion sets the congestion bits in the Frame Relay header to 1. As a result, transmitted traffic has the FECN bit set to 1, and return traffic on the same virtual circuit has the BECN bit set to 1.

When the FECN and BECN bits are set to 1, they provide a congestion notification to the source and destination devices. The devices can respond in either of two ways: to control traffic on the circuit by sending it through other routes, or to reduce the load on the circuit by discarding packets.

If devices discard packets as a means of congestion (flow) control, Frame Relay uses the discard eligibility (DE) bit to give preference to some packets in discard decisions. A DE value of 1 indicates that the frame is of lower importance than other frames and more likely to be dropped during congestion. Critical data (such as signaling protocol messages) without the DE bit set is less likely to be dropped.

### Related Documentation

- [Understanding Physical Encapsulation on an Interface on page 199](#)

## Understanding Point-to-Point Protocol

---

**Supported Platforms** [SRX1500](#), [SRX300](#), [SRX320](#), [SRX340](#)

The Point-to-Point Protocol (PPP) is an encapsulation protocol for transporting IP traffic across point-to-point links. PPP is made up of three primary components:

- Link Control Protocol (LCP)—Establishes working connections between two points.
- Authentication protocol—Enables secure connections between two points.
- Network control protocol (NCP)—Initializes the PPP protocol stack to handle multiple Network Layer protocols, such as IPv4, IPv6, and Connectionless Network Protocol (CLNP).

This topic contains the following sections:

- [Link Control Protocol on page 202](#)
- [PPP Authentication on page 203](#)
- [Network Control Protocols on page 203](#)
- [Magic Numbers on page 204](#)
- [CSU/DSU Devices on page 204](#)

### Link Control Protocol

LCP is responsible for establishing, maintaining, and tearing down a connection between two endpoints. LCP also tests the link and determines whether it is active. LCP establishes a point-to-point connection as follows:

1. LCP must first detect a clocking signal on each endpoint. However, because the clocking signal can be generated by a network clock and shared with devices on the network, the presence of a clocking signal is only a preliminary indication that the link might be functioning.
2. When a clocking signal is detected, a PPP host begins transmitting PPP Configure-Request packets.
3. If the remote endpoint on the point-to-point link receives the Configure-Request packet, it transmits a Configure-Acknowledgement packet to the source of the request.
4. After receiving the acknowledgement, the initiating endpoint identifies the link as established. At the same time, the remote endpoint sends its own request packets and processes the acknowledgement packets. In a functioning network, both endpoints treat the connection as established.

During connection establishment, LCP also negotiates connection parameters such as FCS and HDLC framing. By default, PPP uses a 16-bit FCS, but you can configure PPP to use either a 32-bit FCS or a 0-bit FCS (no FCS). Alternatively, you can enable HDLC encapsulation across the PPP connection.

After a connection is established, PPP hosts generate Echo-Request and Echo-Response packets to maintain a PPP link.

## PPP Authentication

PPP's authentication layer uses a protocol to help ensure that the endpoint of a PPP link is a valid device. Authentication protocols include the Password Authentication Protocol (PAP), the Extensible Authentication Protocol (EAP), and the Challenge Handshake Authentication Protocol (CHAP). CHAP is the most commonly used.



**NOTE:** Support for user id and the password to comply with full ASCII character set is supported through RFC 2486.

The user can enable or disable the RFC 2486 support under the PPP options. The RFC 2486 is disabled by default, and enable the support globally use the command `set access ppp-options compliance rfc 2486`.

CHAP ensures secure connections across PPP links. After a PPP link is established by LCP, the PPP hosts at either end of the link initiate a three-way CHAP handshake. Two separate CHAP handshakes are required before both sides identify the PPP link as established.

CHAP configuration requires each endpoint on a PPP link to use a shared secret (password) to authenticate challenges. The shared secret is never transmitted over the wire. Instead, the hosts on the PPP connection exchange information that enables both to determine that they share the same secret. Challenges consist of a hash function calculated from the secret, a numeric identifier, and a randomly chosen challenge value that changes with each challenge. If the response value matches the challenge value, authentication is successful. Because the secret is never transmitted and is required to calculate the challenge response, CHAP is considered very secure.

PAP authentication protocol uses a simple two-way handshake to establish identity. PAP is used after the link establishment phase (LCP up), during the authentication phase. Junos OS can support PAP in one direction (egress or ingress), and CHAP in the other.

## Network Control Protocols

After authentication is completed, the PPP connection is fully established. At this point, any higher level protocols (for example, IP protocols) can initialize and perform their own negotiations and authentication.

PPP NCPs include support for the following protocols. IPCP and IPv6CP are the most widely used on SRX Series devices.

- IPCP—IP Control Protocol
- IPv6CP—IPv6 Control Protocol
- OSINLCP—OSI Network Layer Control Protocol (includes IS-IS, ES-IS, CLNP, and IDRP)

## Magic Numbers

Hosts running PPP can create “magic” numbers for diagnosing the health of a connection. A PPP host generates a random 32-bit number and sends it to the remote endpoint during LCP negotiation and echo exchanges.

In a typical network, each host's magic number is different. A magic number mismatch in an LCP message informs a host that the connection is not in loopback mode and traffic is being exchanged bidirectionally. If the magic number in the LCP message is the same as the configured magic number, the host determines that the connection is in loopback mode, with traffic looped back to the transmitting host.

Looping traffic back to the originating host is a valuable way to diagnose network health between the host and the loopback location. To enable loopback testing, telecommunications equipment typically supports channel service unit/data service unit (CSU/DSU) devices.

## CSU/DSU Devices

A channel service unit (CSU) connects a terminal to a digital line. A data service unit (DSU) performs protective and diagnostic functions for a telecommunications line. Typically, the two devices are packaged as a single unit. A CSU/DSU device is required for both ends of a T1 or T3 connection, and the units at both ends must be set to the same communications standard.

A CSU/DSU device enables frames sent along a link to be looped back to the originating host. Receipt of the transmitted frames indicates that the link is functioning correctly up to the point of loopback. By configuring CSU/DSU devices to loop back at different points in a connection, network operators can diagnose and troubleshoot individual segments in a circuit.

### Related Documentation

- [Understanding Physical Encapsulation on an Interface on page 199](#)

---

## Understanding High-Level Data Link Control

**Supported Platforms** SRX1500, SRX320, SRX340

High-Level Data Link Control (HDLC) is a bit-oriented, switched and nonswitched link-layer protocol. HDLC is widely used because it supports half-duplex and full-duplex connections, point-to-point and point-to-multipoint networks, and switched and nonswitched channels.

This topic contains the following sections:

- [HDLC Stations on page 205](#)
- [HDLC Operational Modes on page 205](#)

## HDLC Stations

Nodes within a network running HDLC are called stations. HDLC supports three types of stations for data link control:

- Primary stations—Responsible for controlling the secondary and combined other stations on the link. Depending on the HDLC mode, the primary station is responsible for issuing acknowledgement packets to allow data transmission from secondary stations.
- Secondary stations—Controlled by the primary station. Under normal circumstances, secondary stations cannot control data transmission across the link with the primary station, are active only when requested by the primary station, and can respond to the primary station only (not to other secondary stations). All secondary station frames are response frames.
- Combined stations—A combination of primary and secondary stations. On an HDLC link, all combined stations can send and receive commands and responses without any permission from any other stations on the link and cannot be controlled by any other station.

## HDLC Operational Modes

HDLC runs in three separate modes:

- Normal Response Mode (NRM)—The primary station on the HDLC link initiates all information transfers with secondary stations. A secondary station on the link can transmit a response of one or more information frames only when it receives explicit permission from the primary station. When the last frame is transmitted, the secondary station must wait for explicit permission before it can transmit more frames.

NRM is used most widely for point-to-multipoint links, in which a single primary station controls many secondary stations.

- Asynchronous Response Mode (ARM)—The secondary station can transmit either data or control traffic at any time, without explicit permission from the primary station. The primary station is responsible for error recovery and link setup, but the secondary station can transmit information at any time.

ARM is used most commonly with point-to-point links, because it reduces the overhead on the link by eliminating the need for control packets.

- Asynchronous Balance Mode (ABM)—All stations are combined stations. Because no other station can control a combined station, all stations can transmit information without explicit permission from any other station. ABM is not a widely used HDLC mode.

### Related Documentation

- [Understanding Physical Encapsulation on an Interface on page 199](#)



# Configuring Point-to-Point Protocol over Ethernet

- Understanding Point-to-Point Protocol over Ethernet on page 207
- Understanding PPPoE Interfaces on page 210
- Example: Configuring PPPoE Interfaces on page 211
- Understanding PPPoE Ethernet Interfaces on page 217
- Example: Configuring PPPoE Encapsulation on an Ethernet Interface on page 217
- Understanding PPPoE ATM-over-ADSL and ATM-over-SHDSL Interfaces on page 218
- Example: Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface on page 219
- Understanding CHAP Authentication on a PPPoE Interface on page 221
- Example: Configuring CHAP Authentication on a PPPoE Interface on page 221
- Verifying Credit-Flow Control on page 223
- Verifying PPPoE Interfaces on page 224
- Verifying R2CP Interfaces on page 225
- Displaying Statistics for PPPoE on page 226
- Setting Tracing Options for PPPoE on page 226

## Understanding Point-to-Point Protocol over Ethernet

---

**Supported Platforms** SRX1500, SRX300, SRX320, SRX340

*Point-to-Point Protocol over Ethernet* (PPPoE) combines PPP, which typically runs over broadband connections, with the Ethernet link-layer protocol that allows users to connect to a network of hosts over a bridge or access concentrator. PPPoE enables service providers to maintain access control through PPP connections and also manage multiple hosts at a remote site.

PPPoE connects multiple hosts on an Ethernet LAN to a remote site through a single customer premises equipment (CPE) device—a Juniper Networks device. Hosts share a common digital subscriber line (DSL), a cable modem, or a wireless connection to the Internet.

To use PPPoE, you must initiate a PPPoE session, encapsulate Point-to-Point Protocol (PPP) packets over Ethernet, and configure the device as a PPPoE client. To provide a

PPPoE connection, each PPP session must learn the Ethernet address of the remote peer and establish a unique session identifier during the PPPoE discovery and session stages.



**NOTE:** Juniper Networks devices with asymmetric digital subscriber line (ADSL) or symmetric high-speed DSL (SHDSL) interfaces can use PPPoE over Asynchronous Transfer Mode (ATM) to connect through DSL lines only, not for direct ATM connections.

PPPoE has two stages, the discovery stage and the PPPoE session stage. In the *discovery stage*, the client discovers the access concentrator by identifying the Ethernet media access control (MAC) address of the access concentrator and establishing a PPPoE session ID. In the *session stage*, the client and the access concentrator build a point-to-point connection over Ethernet, based on the information collected in the discovery stage.

This topic contains the following sections:

- [PPPoE Discovery Stage on page 208](#)
- [PPPoE Session Stage on page 209](#)

## PPPoE Discovery Stage

To initiate a PPPoE session, a host must first identify the Ethernet MAC address of the remote peer and establish a unique PPPoE session ID for the session. Learning the remote Ethernet MAC address is called *PPPoE discovery*.

During the PPPoE discovery process, the host does not discover a remote endpoint on the Ethernet network. Instead, the host discovers the access concentrator through which all PPPoE sessions are established. Discovery is a client/server relationship, with the host (a device running Junos OS) acting as the client and the access concentrator acting as the server. Because the network might have more than one access concentrator, the discovery stage allows the client to communicate with all of them and select one.



**NOTE:** A device cannot receive PPPoE packets from two different access concentrators on the same physical interface.

The PPPoE discovery stage consists of the following steps:

1. PPPoE Active Discovery Initiation (PADI)—The client initiates a session by broadcasting a PADI packet to the LAN to request a service.
2. PPPoE Active Discovery Offer (PADO)—Any access concentrator that can provide the service requested by the client in the PADI packet replies with a PADO packet that contains its own name, the unicast address of the client, and the service requested. An access concentrator can also use the PADO packet to offer other services to the client.

3. PPPoE Active Discovery Request (PADR)—From the PADOs it receives, the client selects one access concentrator based on its name or the services offered and sends it a PADR packet to indicate the service or services needed.
4. PPPoE Active Discovery Session-Confirmation (PADS)—When the selected access concentrator receives the PADR packet, it accepts or rejects the PPPoE session:
  - To accept the session, the access concentrator sends the client a PADS packet with a unique session ID for a PPPoE session and a service name that identifies the service under which it accepts the session.
  - To reject the session, the access concentrator sends the client a PADS packet with a service name error and resets the session ID to zero.

## PPPoE Session Stage

The PPPoE session stage starts after the PPPoE discovery stage is over. The access concentrator can start the PPPoE session after it sends a PADS packet to the client, or the client can start the PPPoE session after it receives a PADS packet from the access concentrator. A device supports multiple PPPoE sessions on each interface, but no more than 256 PPPoE sessions per device.

Each PPPoE session is uniquely identified by the Ethernet address of the peer and the session ID. After the PPPoE session is established, data is sent as in any other PPP encapsulation. The PPPoE information is encapsulated within an Ethernet frame and is sent to a unicast address. Magic numbers, echo requests, and all other PPP traffic behave exactly as in normal PPP sessions. In this stage, both the client and the server must allocate resources for the PPPoE logical interface.

After a session is established, the client or the access concentrator can send a PPPoE Active Discovery Termination (PADT) packet anytime to terminate the session. The PADT packet contains the destination address of the peer and the session ID of the session to be terminated. After this packet is sent, the session is closed to PPPoE traffic.



**NOTE:** If PPPoE session is already up and the user restarts the PPPoE daemon, a new PPPoE daemon with a new PID starts while the existing session is not terminated.

If PPPoE session is already down and user restarts the PPPoE daemon, the PPPoE discovery establishes a new session.

The PPPoE session is not terminated for the following configuration changes:

- Changing idle time out value
- Changing auto rec timer value
- Deleting idle time out
- Deleting auto rec timer
- Add new auto rec time

- Add new idle time out
- Change negotiate address to static address
- Change static ip address to a new static ip address
- Changing default chap secrete

The PPPoE session is terminated for the following configuration changes:

- Add ac name
- Delete chap ppp options
- Add new chap ppp options
- Configure uifd mac



**NOTE:** When the MTU for an underlying physical interface is changed, it brings down the PPPoE session. For PPPoE, an MTU greater than 1492 cannot be achieved.

**Related Documentation**

- [Understanding Physical Encapsulation on an Interface on page 199](#)
- [Understanding PPPoE Interfaces on page 210](#)
- [Understanding PPPoE Ethernet Interfaces on page 217](#)
- [Understanding PPPoE ATM-over-ADSL and ATM-over-SHDSL Interfaces on page 218](#)
- [Understanding CHAP Authentication on a PPPoE Interface on page 221](#)
- [Understanding the PPPoE-Based Radio-to-Router Protocol on page 230](#)

## Understanding PPPoE Interfaces

---

**Supported Platforms** [SRX1500](#), [SRX300](#), [SRX320](#), [SRX340](#)

The device's Point-to-Point Protocol over Ethernet (PPPoE) interface to the access concentrator can be a Fast Ethernet interface, a Gigabit Ethernet interface, a redundant Ethernet interface, an ATM-over-ADSL interface, or an ATM-over-SHDSL interface. The PPPoE configuration is the same for all interfaces. The only difference is the encapsulation for the underlying interface to the access concentrator:

- If the interface is Ethernet, use a PPPoE encapsulation.
- If the interface is ATM-over-ADSL or ATM-over-SHDSL, use a PPPoE over ATM encapsulation.

To configure a PPPoE interface, you create an interface with a logical interface unit 0, then specify a logical Ethernet or ATM interface as the underlying interface for the PPPoE session. You then specify other PPPoE options, including the access concentrator and PPPoE session parameters.



**NOTE:** PPPoE over redundant Ethernet (reth) interface is supported on SRX300, SRX320, and SRX340 devices. This feature allows an existing PPPoE session to continue without starting a new PPPoE session in the event of a failover.

#### Related Documentation

- [Understanding Point-to-Point Protocol on page 202](#)
- [Example: Configuring PPPoE Interfaces on page 211](#)

## Example: Configuring PPPoE Interfaces

**Supported Platforms** SRX1500, SRX300, SRX320, SRX340, SRX550M

This example shows how to configure a PPPoE interface.

- [Requirements on page 211](#)
- [Overview on page 211](#)
- [Configuration on page 211](#)
- [Disabling the End-of-List Tag on page 215](#)

### Requirements

Before you begin, configure an Ethernet interface. See “[Example: Creating an Ethernet Interface](#)” on page 122.

### Overview

In this example, you create the PPPoE interface pp0.0 and specify the logical Ethernet interface ge-0/0/1.0 as the underlying interface. You also set the access concentrator, set the PPPoE session parameters, and set the MTU of the IPv4 family to 1492.

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces pp0 unit 0 pppoe-options underlying-interface ge-0/0/1.0
 access-concentrator ispl.com auto-reconnect 100 idle-timeout 100 client service-name
 video@ispl.com
set interfaces pp0 unit 0 family inet mtu 1492 negotiate-address
```

#### Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a PPPoE interface:

1. Create a PPPoE interface.

- ```
[edit]
user@host# edit interfaces pp0 unit 0
```
- Configure PPPoE options.

```
[edit interfaces pp0 unit 0]
user@host# set pppoe-options underlying-interface ge-0/0/1.0 access-concentrator
ispl.com auto-reconnect 100 idle-timeout 100 client service-name video@ispl.com
```
 - Configure the MTU.

```
[edit interfaces pp0 unit 0]
user@host# set family inet mtu 1492
```
 - Configure the PPPoE interface address.

```
[edit interfaces pp0 unit 0]
user@host# set family inet negotiate-address
```

Results From configuration mode, confirm your configuration by entering the **show interfaces pp0** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces pp0
unit 0 {
  pppoe-options {
    underlying-interface ge-0/0/1.0;
    idle-timeout 100;
    access-concentrator ispl.com;
    service-name "vide0@ispl.com";
    auto-reconnect 100;
    client;
  }
  family inet {
    mtu 1492;
    negotiate-address;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying PPPoE Interfaces on page 212](#)
- [Verifying PPPoE Sessions on page 213](#)
- [Verifying the PPPoE Version on page 214](#)
- [Verifying PPPoE Statistics on page 214](#)

Verifying PPPoE Interfaces

Purpose Verify that the PPPoE device interfaces are configured properly.

Action From operational mode, enter the **show interfaces pp0** command.

```

user@host> show interfaces pp0
Physical interface: pp0, Enabled, Physical link is Up
  Interface index: 67, SNMP ifIndex: 317
  Type: PPPoE, Link-level type: PPPoE, MTU: 9192
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Last flapped   : Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

Logical interface pp0.0 (Index 1) (SNMP ifIndex 330)
  Flags: Point-To-Point SNMP-Traps 16384 Encapsulation: PPPoE
  PPPoE:
    State: SessionUp, Session ID: 3304,
    Session AC name: ispl.com, AC MAC address: 00:90:1a:40:f6:4c,
    Service name: video@ispl.com, Configured AC name: ispl.com,
    Auto-reconnect timeout: 60 seconds
    Underlying interface: ge-5/0/0.0 (Index 71)
  Input packets : 23
  Output packets: 22
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 16 (00:00:26 ago), Output: 0 (never)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mp1s:
  Not-configured
  CHAP state: Success
    Protocol inet, MTU: 1492
    Flags: Negotiate-Address
    Addresses, Flags: Kernel Is-Preferred Is-Primary
    Destination: 211.211.211.2, Local: 211.211.211.1

```

The output shows information about the physical and the logical interfaces. Verify the following information:

- The physical interface is enabled and the link is up.
- The PPPoE session is running on the correct logical interface.
- For state, the state is active (up).
- For underlying interface, the physical interface on which the PPPoE session is running is correct:
 - For an Ethernet connection, the underlying interface is Fast Ethernet or Gigabit Ethernet—for example, ge-5/0/0.0.
 - For an ATM-over-ADSL or ATM-over-SHDSL connection, the underlying interface is ATM—for example, at-2/0/0.0.

Verifying PPPoE Sessions

Purpose Verify that a PPPoE session is running properly on the logical interface.

Action From operational mode, enter the **show pppoe interfaces** command.

```
user@host> show pppoe interfaces
pp0.0 Index 67
  State: Session up, Session ID: 31,
  Service name: video@isp1.com, Configured AC name: isp1.com,
  Session AC name: belur, AC MAC address: 00:90:1a:40:f6:4e,
  Auto-reconnect timeout: 1 seconds,
  Underlying interface: ge-0/0/1.0 Index 69
```

The output shows information about the PPPoE sessions. Verify the following information:

- The PPPoE session is running on the correct logical interface.
- For state, the session is active (up).
- For underlying interface, the physical interface on which the PPPoE session is running is correct:
 - For an Ethernet connection, the underlying interface is Fast Ethernet or Gigabit Ethernet—for example, ge-0/0/1.0.
 - For an ATM-over-ADSL or ATM-over-SHDSL connection, the underlying interface is ATM—for example, at-2/0/0.0.



NOTE: To clear a PPPoE session on the pp0.0 interface, use the **clear pppoe sessions pp0.0** command. To clear all sessions on the interface, use the **clear pppoe sessions** command.

Verifying the PPPoE Version

Purpose Verify the version information of the PPPoE protocol configured on the device interfaces.

Action From operational mode, enter the **show pppoe version** command.

```
user@host> show pppoe version
Point-to-Point Protocol Over Ethernet, version 1. rfc2516
  PPPoE protocol           = Enabled
  Maximum Sessions         = 256
  PADI resend timeout      = 2 seconds
  PADR resend timeout      = 16 seconds
  Max resend timeout       = 64 seconds
  Max Configured AC timeout = 4 seconds
```

The output shows PPPoE protocol information. Verify the following information:

- The correct version of the PPPoE protocol is configured on the interface.
- For PPPoE protocol, the PPPoE protocol is enabled.

Verifying PPPoE Statistics

Purpose Verify the statistics information about PPPoE interfaces.

Action From operational mode, enter the **show pppoe statistics** command.

```
user@host> show pppoe statistics
Active PPPoE sessions: 4
  PacketType          Sent      Received
  PADI                502       0
  PADO                 0        219
  PADR                219       0
  PADS                 0        219
  PADT                 0        161
  Service name error  0         0
  AC system error     0         13
  Generic error       0         0
  Malformed packets  0         41
  Unknown packets    0         0
Timeout
  PADI                42
  PADO                 0
  PADR                 0
```

The output shows information about active sessions on PPPoE interfaces. Verify the following information:

- Total number of active PPPoE sessions running on the interface
- For packet type, the number of packets of each type sent and received during the PPPoE session

Disabling the End-of-List Tag

During the PPPoE discovery stage, any access concentrator that can provide the service requested by the client in the PADI packet replies with a PADO packet that contains its own name, the unicast address of the client, and the service requested. An access concentrator can also use the PADO packet to offer other services to the client. When a client receives a PADO packet, and if it encounters the **End-of-List** tag in the PADO packet, tags after the **End-of-List** tag are ignored and the complete information is not processed correctly. As a result, the PPPoE connection is not established correctly.

You can configure the **ignore-eol-tag** option to disable the **End-of-List** tag in the PADO packet.

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To disable the **End-of-List** tag:

1. Create a PPPoE interface.

```
[edit]
user@host# set interfaces pp0 unit 0
```

2. Configure PPPoE options.

```
[edit interfaces pp0 unit 0]
user@host# set pppoe-options ignore-eol-tag
```

Results From configuration mode, confirm your configuration by entering the **show interfaces pp0** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces pp0
unit 0 {
  pppoe-options {
    ignore-eol-tag;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verifying That the End-of-List Tag Is Disabled

Purpose Verify the status of the **End-of-List** tag in the PPPoE configuration.

Action From operational mode, enter the **show interfaces pp0.0** command.

```
user@host> show pppoe interfaces pp0.0
Logical interface pp0.0 (Index 78) (SNMP ifIndex 541)
  Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: PPPoE
  PPPoE:
    State: SessionUp, Session ID: 3,
    Session AC name: cell, Remote MAC address: 00:26:88:f7:77:83,
    Configured AC name: None, Service name: None,
    Auto-reconnect timeout: Never, Idle timeout: Never,
    Underlying interface: ge-0/0/3.0 (Index 77)
  Ignore End-Of-List tag: Enable
```

```
user@host> show pppoe interfaces pp0.0 extensive
pp0.0 Index 74
  State: Session up, Session ID: 1,
  Service name: None,
  Session AC name: cell, Configured AC name: None,
  Remote MAC address: 00:26:88:f7:77:83,
  Session uptime: 00:02:03 ago,
  Auto-reconnect timeout: 10 seconds, Idle timeout: Never,
  Underlying interface: ge-0/0/3.0 Index 73
  Ignore End-of-List tag: Enable
  PacketType          Sent      Received
  PADI                 23         0
  PADO                 0          5
  PADR                 11         0
  PADS                 0          2
  PADT                 2          0
  Service name error   0          0
  AC system error      0          0
  Generic error        0          0
  Malformed packets    0          0
  Unknown packets      0          0
  Timeout
  PADI                 3          0
  PADO                 0          0
  PADR                 3          0
  Receive Error Counters
  PADI                 0          0
  PADO                 0          0
```

```

PADR          0
PADS          0

```

The output shows information about active sessions on PPPoE interfaces. Verify that the **Ignore End-of-List tag: Enable** option is set.

Related Documentation

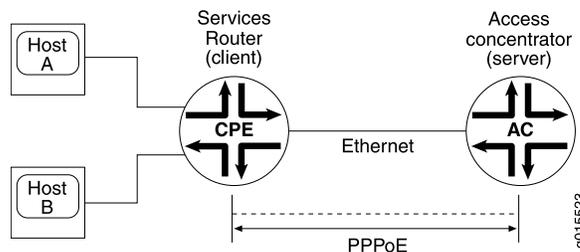
- [Understanding PPPoE Interfaces on page 210](#)

Understanding PPPoE Ethernet Interfaces

Supported Platforms SRX1500, SRX300, SRX320, SRX340

During a Point-to-Point Protocol over Ethernet (PPPoE) session, the device encapsulates each PPP frame in an Ethernet frame and transports the frames over an Ethernet loop. [Figure 13 on page 217](#) shows a typical PPPoE session between a device and an access concentrator on the Ethernet loop.

Figure 13: PPPoE Session on the Ethernet Loop



To configure PPPoE on an Ethernet interface, you configure encapsulation on the logical interface.

Related Documentation

- [Understanding Point-to-Point Protocol over Ethernet on page 207](#)
- [Example: Configuring PPPoE Encapsulation on an Ethernet Interface on page 217](#)

Example: Configuring PPPoE Encapsulation on an Ethernet Interface

Supported Platforms SRX1500, SRX300, SRX320, SRX340

This example shows how to configure PPPoE encapsulation on an Ethernet interface.

- [Requirements on page 217](#)
- [Overview on page 218](#)
- [Configuration on page 218](#)
- [Verification on page 218](#)

Requirements

Before you begin:

- Configure an Ethernet interface. See [“Example: Creating an Ethernet Interface”](#) on page 122.
- Configure a PPPoE encapsulation interface. See [“Example: Configuring PPPoE Interfaces”](#) on page 211.

Overview

In this example, you configure PPPoE encapsulation on the ge-0/0/1 interface.

Configuration

Step-by-Step Procedure

To configure PPPoE encapsulation:

1. Enable PPPoE encapsulation on the interface.

[edit]

```
user@host# set interfaces ge-0/0/1 unit 0 encapsulation ppp-over-ether
```

2. Commit the configuration if you are done configuring the device.

[edit]

```
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show interfaces ge-0/0/1** command.

Related Documentation

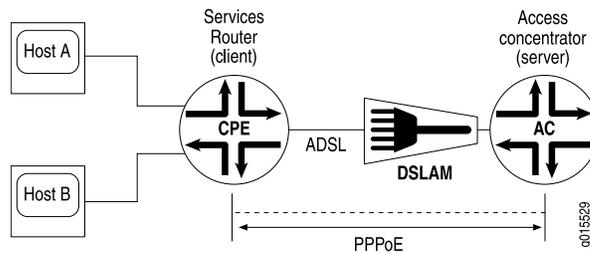
- [Understanding PPPoE Ethernet Interfaces on page 217](#)

Understanding PPPoE ATM-over-ADSL and ATM-over-SHDSL Interfaces

Supported Platforms [SRX210, SRX220, SRX240](#)

When an ATM network is configured with a point-to-point connection, Point-to-Point Protocol over Ethernet (PPPoE) can use ATM Adaptation Layer 5 (AAL5) for framing PPPoE-encapsulated packets. The AAL5 protocol provides a virtual connection between the client and the server within the same network. The device encapsulates each PPPoE frame in an ATM frame and transports each frame over an asymmetric digital subscriber line (ADSL) or symmetric high-speed DSL (SHDSL) loop and a digital subscriber line access multiplexer (DSLAM). For example, [Figure 14 on page 219](#) shows a typical PPPoE over ATM session between a device and an access concentrator on an ADSL loop.

Figure 14: PPPoE Session on an ADSL Loop



For PPPoE on an ATM-over-ADSL or ATM-over-SHDSL interface, you must configure encapsulation on both the physical and logical interfaces. To configure encapsulation on an ATM-over-ADSL or ATM-over-SHDSL physical interface, use Ethernet over ATM encapsulation. To configure encapsulation on an ATM-over-ADSL or ATM-over-SHDSL logical interface, use PPPoE over AAL5 logical link control (LLC) encapsulation. LLC encapsulation allows a single ATM virtual connection to transport multiple protocols.

Related Documentation

- [Understanding Point-to-Point Protocol over Ethernet on page 207](#)
- [Example: Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface on page 219](#)

Example: Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface

Supported Platforms SRX210, SRX220, SRX240

This example shows how to configure a physical interface for Ethernet over ATM encapsulation and how to create a logical interface for PPPoE over LLC encapsulation.

- [Requirements on page 219](#)
- [Overview on page 219](#)
- [Configuration on page 220](#)
- [Verification on page 221](#)

Requirements

Before you begin:

- Configure network interfaces. See [“Example: Creating an Ethernet Interface” on page 122](#).
- Configure PPPoE interfaces. See [“Example: Configuring PPPoE Interfaces” on page 211](#).
- Configure PPPoE encapsulation on an Ethernet interface. See [“Example: Configuring PPPoE Encapsulation on an Ethernet Interface” on page 217](#).

Overview

In this example, you configure the physical interface at-2/0/0 for Ethernet over ATM encapsulation. As part of the configuration, you set the virtual path identifier (VPI) on an ATM-over-ADSL physical interface to 0, you set the ADSL operating mode to auto, and you set the encapsulation type to ATM-over-ADSL. Then you create a logical interface for PPPoE over LLC encapsulation.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces at-2/0/0 atm-options vpi 0
set interfaces at-2/0/0 dsl-options operating-mode auto
set interfaces at-2/0/0 encapsulation ethernet-over-atm
set interfaces at-2/0/0 unit 0 encapsulation ppp-over-ether-over-atm-llc vci 0.120
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure PPPoE encapsulation on an ATM-over-ADSL interface:

1. Configure the physical interface.

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

2. Set the VPI on the interface.

```
[edit interfaces at-2/0/0]
user@host# set atm-options vpi 0
```

3. Configure the ADSL operating mode.

```
[edit interfaces at-2/0/0]
user@host# set dsl-options operating-mode auto
```

4. Configure PPPoE encapsulation.

```
[edit interfaces at-2/0/0]
user@host# set encapsulation ethernet-over-atm
```

5. Create a logical interface and configure LLC encapsulation.

```
[edit interfaces at-2/0/0]
user@host# set unit 0 encapsulation ppp-over-ether-over-atm-llc vci 0.120
```

Results From configuration mode, confirm your configuration by entering the **show interfaces at-2/0/0** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces at-2/0/0 {
  encapsulation ethernet-over-atm;
  atm-options {
    vpi 0;
  }
  dsl-options {
    operating-mode auto;
  }
  unit 0 {
```

```

    encapsulation ppp-over-ether-over-atm-llc;
    vci 0.120;
  }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying a PPPoE Configuration for an ATM-over-ADSL or ATM-over-SHDSL Interface on page 221](#)

[Verifying a PPPoE Configuration for an ATM-over-ADSL or ATM-over-SHDSL Interface](#)

Purpose Verify the PPPoE configuration for an ATM-over-ADSL or ATM-over-SHDSL interface.

Action From operational mode, enter the **show interfaces** command.

Related Documentation

- [Understanding PPPoE ATM-over-ADSL and ATM-over-SHDSL Interfaces on page 218](#)

[Understanding CHAP Authentication on a PPPoE Interface](#)

Supported Platforms [SRX1500, SRX300, SRX320, SRX340](#)

For interfaces with Point-to-Point Protocol over Ethernet (PPPoE) encapsulation, you can configure interfaces to support the PPP Challenge Handshake Authentication Protocol (CHAP). When you enable CHAP on an interface, the interface can authenticate its peer and be authenticated by its peer.

If you set the **passive** option to handle incoming CHAP packets only, the interface does not challenge its peer. However, if the interface is challenged, it responds to the challenge. If you do not set the **passive** option, the interface always challenges its peer.

You can configure Remote Authentication Dial-In User Service (RADIUS) authentication of PPP sessions using CHAP. CHAP enables you to send RADIUS messages through a routing instance to customer RADIUS servers in a private network.

Related Documentation

- [Understanding Point-to-Point Protocol over Ethernet on page 207](#)
- [Example: Configuring CHAP Authentication on a PPPoE Interface on page 221](#)

[Example: Configuring CHAP Authentication on a PPPoE Interface](#)

Supported Platforms [SRX1500, SRX300, SRX320, SRX340](#)

This example shows how to configure CHAP authentication on a PPPoE interface.

- [Requirements on page 222](#)
- [Overview on page 222](#)
- [Configuration on page 222](#)
- [Verification on page 223](#)

Requirements

Before you begin:

- Configure an Ethernet interface. See “[Example: Creating an Ethernet Interface](#)” on [page 122](#).
- Configure a PPPoE interface. See “[Example: Configuring PPPoE Interfaces](#)” on [page 211](#).
- Configure PPPoE encapsulation on an ATM-over-ADSL interface. See “[Example: Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface](#)” on [page 219](#).

Overview

In this example, you configure a CHAP access profile, and then apply it to the PPPoE interface pp0. You also configure the hostname to be used in CHAP challenge and response packets, and set the passive option for handling incoming CHAP packets.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set access profile A-ppp-client client client1 chap-secret my-secret
set interfaces pp0 unit 0 ppp-options chap access-profile A-ppp-client local-name
A-ge-0/0/1.0 passive
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure CHAP on a PPPoE interface:

1. Configure a CHAP access profile.

```
[edit]
user@host# set access profile A-ppp-client client client1 chap-secret my-secret
```
2. Enable CHAP options on the interface.

```
[edit]
user@host# edit interfaces pp0 unit 0 ppp-options chap
```
3. Configure the CHAP access profile on the interface.

```
[edit interfaces pp0 unit 0 ppp-options chap]
```

```
user@host# set access-profile A-ppp-client
```

4. Configure a hostname for the CHAP challenge and response packets.

```
[edit interfaces pp0 unit 0 ppp-options chap]
user@host# set local-name A-ge-0/0/1.0
```

5. Set the passive option to handle incoming CHAP packets only.

```
[edit interfaces pp0 unit 0 ppp-options chap]
user@host# set passive
```

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

```
[edit]
user@host# show interfaces
pp0 {
  unit 0 {
    ppp-options {
      chap {
        access-profile A-ppp-client;
        local-name A-ge-0/0/1.0;
        passive;
      }
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Verifying CHAP Authentication

Purpose Verify that CHAP is enabled on the interface.

Action From operational mode, enter the **show interfaces** command.

Related Documentation

- [Understanding CHAP Authentication on a PPPoE Interface on page 221](#)

Verifying Credit-Flow Control

Purpose Display PPPoE credit-flow control information about credits on each side of the PPPoE session when credit processing is enabled on the interface.

Action user@host> **show pppoe interface detail**

```
pp0.51 Index 73
State: Session up, Session ID: 3,
Service name: None,
```

```
Configured AC name: None, Session AC name: None,
Remote MAC address: 00:22:83:84:2e:81,
Session uptime: 00:05:48 ago,
Auto-reconnect timeout: Never, Idle timeout: Never,
Underlying interface: ge-0/0/4.1 Index 72
PADG Credits: Local: 12345, Remote: 6789, Scale factor: 128 bytes
PADQ Current bandwidth: 750 Kbps, Maximum 1000 Kbps
Quality: 85, Resources 65, Latency 100 msec.
Dynamic bandwidth: 3 Kbps
```

```
pp0.1000 Index 71
State: Down, Session ID: 1,
Service name: None,
Configured AC name: None, Session AC name: None,
Remote MAC address: 00:00:00:00:00:00,
Auto-reconnect timeout: Never, Idle timeout: Never,
Underlying interface: ge-0/0/1.0 Index 70
PADG Credits: enabled
Dynamic bandwidth: enabled
```

- Related Documentation**
- [Understanding CHAP Authentication on a PPPoE Interface on page 221](#)
 - [Verifying Credit-Flow Control on page 223](#)

Verifying PPPoE Interfaces

Purpose Display PPPoE interfaces information.

- Action**
- To display PPPoE interface information:

```
user@host> show pppoe interfaces pp0.51 detail
```

```
pp0.51 Index 75
State: Session up, Session ID: 1,
Service name: None,
Configured AC name: None, Session AC name: None,
Remote MAC address: 00:11:22:33:44:55,
Session uptime: 00:04:18 ago,
Auto-reconnect timeout: Never, Idle timeout: Never,
Underlying interface: ge-0/0/1.0 Index 70
PADQ Current bandwidth: 750 Kbps, Maximum 1000 Kbps
Quality: 85, Resources 65, Latency 100 msec.
Dynamic bandwidth: 3 Kbps
```

- To display PPPoE terse interface information:

```
user@host> show pppoe interfaces terse pp0.51
```

```
Interface  Admin Link Proto  Local          Remote
pp0.51    up up  inet  5.1.1.1      --> 5.1.1.2
          inet6   fe80::21f:12ff:fed2:2918/64
                               feee::5:1:1:1/126
```

- Related Documentation**
- [Understanding PPPoE Interfaces on page 210](#)
 - [Example: Configuring PPPoE Interfaces on page 211](#)

Verifying R2CP Interfaces

Purpose Display R2CP interfaces information.

- Action**
- To display R2CP interface information:

```
root@host> show r2cp interfaces
```

```
Interface: ge-0/0/3.51
Nodes: 0
```

- To display R2CP information:

```
root@host> show r2cp radio extensive
```

Node Packet Type	Sent	Received	Errors
MIM	-	1	0
ROM	1	-	-
Heartbeats	0	0	0
Node Term	0	0	0
Node Term Ack	0	0	-
Heartbeat Timeouts	0		
Node Term Timeouts	0		

Session Packet Type	Sent	Received	Errors
Init	-	1	0
Init ACK	1	-	-
Update	-	0	0
Terminate	0	0	0
Terminate ACK	0	0	0
Terminate Timeouts	0		

- To display R2CP session information:

```
root@host> show r2cp sessions extensive
```

```
Session: 1
Destination MAC address 01:02:03:04:05:06
Status: Established VLANs 201
Virtual channel: 2
Session Update: last received: 3.268 seconds
Current bandwidth: 22000 Kbps, Maximum 22000 Kbps
Quality: 100, Resources 100, Latency 100 msec.
Effective bandwidth: 952 Kbps, last change: 51.484 seconds
Updates below threshold: 1
```

Session Packet Type	Sent	Received	Errors
Init	-	1	0
Init ACK	1	-	-
Update	-	0	0
Terminate	0	0	0
Terminate ACK	0	0	0
Terminate Timeouts	0		

- Related Documentation**
- [Understanding PPPoE Interfaces on page 210](#)
 - [Example: Configuring PPPoE Interfaces on page 211](#)

Displaying Statistics for PPPoE

Purpose Display PPPoE statistics.

Action user@host> show interfaces pp0.51 statistics

```
Logical interface pp0.51 (Index 75) (SNMP ifIndex 137)
  Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: PPPoE
  PPPoE:
    State: SessionUp, Session ID: 1,
    Session AC name: None, Remote MAC address: 00:22:83:84:2f:03,
    Underlying interface: ge-0/0/4.1 (Index 74)
  Input packets : 20865
  Output packets: 284636
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 0 (never), Output: 943 (00:00:06 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Opened, iso: Not-configured, mp1s:
  Not-configured
  CHAP state: Closed
  PAP state: Closed
  Security: Zone: Null
  Protocol inet, MTU: 1492
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
    Destination: 5.1.1.2, Local: 5.1.1.1
  Protocol inet6, MTU: 1492
    Flags: None
    Addresses, Flags: Is-Preferred
    Destination: fe80::21f:12ff:fed2:2918
    Addresses, Flags: Is-Preferred Is-Primary
    Destination: feee::5:1:1:0/126, Local: feee::5:1:1:1
```

- Related Documentation**
- [Understanding CHAP Authentication on a PPPoE Interface on page 221](#)
 - [Verifying Credit-Flow Control on page 223](#)

Setting Tracing Options for PPPoE

To trace the operations of the router's PPPoE process, include the traceoptions statement at the [edit protocols pppoe] hierarchy level:

```
[edit protocols pppoe]
traceoptions {
  file filename <files number> <match regular-expression> <size size> <world-readable |
  no-world-readable>;
  flag flag;
  level severity-level;
  no-remote-trace;
}
```

To specify more than one tracing operation, include multiple **flag** statements.

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

- **all**—All areas of code
- **config**—Configuration code
- **events**—Event code
- **gres**—Gres code
- **init**—Initialization code
- **interface-db**—Interface database code
- **memory**—Memory management code
- **protocol**—PPPoE protocol processing code
- **rtsock**—Routing socket code
- **session-db**—Session management code
- **signal**—Signal handling code
- **state**—State handling code
- **timer**—Timer code
- **ui**—User interface code

**Related
Documentation**

- [Understanding PPPoE Interfaces on page 210](#)
- [Example: Configuring PPPoE Interfaces on page 211](#)

Configuring PPPoE-Based Radio-to-Router Protocol

- [PPPoE-Based Radio-to-Router Protocols Overview on page 229](#)
- [Understanding the PPPoE-Based Radio-to-Router Protocol on page 230](#)
- [Configuring PPPoE-Based Radio-to-Router Protocols on page 232](#)
- [Example: Configuring the PPPoE-Based Radio-to-Router Protocol on page 232](#)
- [Credit Flow Control for PPPoE on page 235](#)
- [PPPoE Credit-Based Flow Control Configuration on page 235](#)

PPPoE-Based Radio-to-Router Protocols Overview

Support for PPPoE-based radio-to-router protocols includes the following extensions to the PPPoE protocol:

- Messages that define how an external device provides the router with timely information about the quality of a link connection
- A flow control mechanism that indicates how much data the router can forward

The router uses the information provided in these PPPoE messages to dynamically adjust the interface speed. When OSPF is notified of this change, it adjusts the cost of the link and updates the routing tables accordingly.

The radio provides ground-to-ground or ground-to-air communications with like devices. When the radio picks up a signal from another device, it initiates a PPPoE session with a directly connected router. The PPPoE session encapsulates the packets that are relayed over a PPP link between the local and remote routers. The remote radio then forwards traffic over an independent PPPoE session between the remote radio and the router to which it is connected. The two routers exchange LCP and IPCP messages to configure the link and exchange OSPF messages to establish the network topology.

The router and radio are deployed in highly dynamic environments, such as moving vehicles. The quality of the radio link between the routers can vary significantly as a vehicle moves behind an obstruction. Each radio monitors the link every 50 milliseconds for changes in the link bandwidth, quality, and utilization. If any changes are detected, the radios announce the new set of metrics to the respective routers through a PPPoE Active Discovery Quality (PADQ) message, which is a nonstandard extension to the

PPPoE Discovery Protocol [RFC2516]. The router transforms these metrics into a bandwidth value for the PPP link and compares it to the value currently in use. When the router detects that the difference exceeds a user-specified threshold, it adjusts the speed of the PPP link. An event message notifies OSPF of the change, which then triggers OSPF to announce any resulting routing topology changes to its neighbors.

The PPPoE-based radio-to-router protocol notifies the router about neighbors joining or leaving the network and to create and maintain OSPF adjacencies over the dynamic links established between them. The costs assigned to these links are based on network conditions and flow control information sent by the radios. The calculations and requests to update interface speeds are performed by routines in a common library.

When PPPoE is used for applications, such as mobile radio, the radio links have variable bandwidth. So a mobile radio can function in a PPPoE environment, PPPoE messaging includes PADQ messages, which enable a link cost to be propagated to OSPF through the evaluation of various link quality metrics. The router uses information from these notifications along with user-configured parameters to calculate interface link costs that are used by the routing protocols.

A radio can send an optional PADQ at any time to query or report link quality metrics. When transmitting PPP streams over radio links, the quality of the link directly affects the throughput. The PADQ packet is used by the radio modem to report link metrics.

To support the credit-based flow control extensions described in RFC4938, PPPoE peers can also grant each other forwarding credits. The grantee can forward traffic to the peer only when it has a sufficient number of credits to do so. Credit-based forwarding allows both sides of the session to agree to use a non-default credit scaling factor during the PADR and PADS message exchange. Although this is used on both sides of the session, this feature provides the radio client with a flow control mechanism that throttles traffic by limiting the number of credits it grants to the router.

**Related
Documentation**

- [Understanding the PPPoE-Based Radio-to-Router Protocol on page 230](#)

Understanding the PPPoE-Based Radio-to-Router Protocol

Supported Platforms [SRX Series](#)

Point-to-Point Protocol over Ethernet (PPPoE)-based radio-to-router protocols include messages that define how an external system will provide the device with timely information about the quality of a link's connection. They also include a flow control mechanism to indicate how much data the device can forward. The device can then use the information provided in the PPPoE messages to dynamically adjust the interface speed of PPP links.

For example, a high-band networking waveform (HNW) radio provides ground-to-ground or ground-to-air communications with like devices. When the HNW picks up a signal from another device, it initiates a PPPoE session with a directly connected device (router). The PPPoE session encapsulates the packets that are relayed over a PPP link between the local and remote devices. The remote radio then forwards traffic to a remote device using an independent PPPoE session. The two devices exchange Link Control Protocol

(LCP) and Internet Protocol Control Protocol (IPCP) messages to configure the link and exchange OSPF messages to establish the network topology.

Each HNW radio monitors the link every 50 milliseconds for changes in the link bandwidth, quality, and utilization. If any changes are detected, the radios announce the new set of metrics to the respective devices through a PPPoE Active Discovery Quality (PADQ) message, which is a nonstandard extension to the PPPoE Discovery Protocol (RFC 2516). The device transforms these metrics into a bandwidth value for the PPP link and compares it to the value currently in use. When the device detects that the difference exceeds a user-specified threshold, it adjusts the speed of the PPP link. OSPF is notified of the change and announces any resulting routing topology changes to its neighbors.

The CLI statement, **radio-router**, indicates that metrics announcements received on the interface will be processed by the device. When a PPPoE logical interface refers to this as an underlying interface, the device then processes incoming PADQ messages and uses information from the host's messages to control the flow of traffic and manage the speed of the link, resulting in a corresponding adjustment of the OSPF cost. If this option is not specified, then PADQ messages received over the underlying interface are ignored.

The following options are available within the **radio-router** configuration statement:

- **bandwidth, resource, latency, and quality** —These statements provide control over the weights used when transforming PADQ link metrics into an interface speed for the virtual link:
 - **bandwidth**—Weight of current (vs. maximum) data rate
 - **resource**—Resource weight
 - **latency**—Latency weight
 - **quality**—Relative link quality weight

All four weights accept values from 0 through 100. The default value for all four weights is 100.

- **credit**—This statement supports the credit-based flow control extensions described in RFC 4938. The statement enables PPPoE peers to grant each other forwarding credits. The grantee is then allowed to forward traffic to the peer only when it has a sufficient number of credits to do so. The subsequent credit interval statement controls how frequently the device generates credit announcement messages. The **interval** sub-statement, which controls the grant rate interval, accepts values from 1 through 60 seconds.
- **threshold**—This statement specifies how much of a difference is required between the calculated and the current interface speeds. The **threshold** value, expressed as a percentage, defaults to 10.

The following hierarchy provides another view of the **radio-router** configuration statements.

```

interfaces{
  interface-name {
    radio-router {
      bandwidth;
    }
  }
}

```

```
        credit {
            interval;
        }
        latency;
        quality;
        resource;
        threshold;
    }
}
```

Related Documentation

- [Understanding Point-to-Point Protocol over Ethernet on page 207](#)
- [Example: Configuring the PPPoE-Based Radio-to-Router Protocol on page 232](#)

Configuring PPPoE-Based Radio-to-Router Protocols

Supported Platforms [SRX Series](#)

To configure the PPPoE-based radio-to-router protocol:

1. Configure PPPoE encapsulation for an Ethernet interface.
2. Configure radio-router on the logical Ethernet interface.
3. Specify the logical Ethernet interface as the underlying interface for the PPPoE session.
4. Configure the operational mode as server.
5. (Optional) Identify the access concentrator by a unique name.
6. Specify how many seconds to wait before attempting to reconnect.
7. Provide a name for the type of service provided by the access concentrator.
8. Configure the maximum transmission unit (MTU) of the interface.
9. Configure the MTU size for the protocol family.
10. Disable the sending of keepalive messages on the logical interface.

Related Documentation

- [Understanding the PPPoE-Based Radio-to-Router Protocol on page 230](#)
- [Example: Configuring the PPPoE-Based Radio-to-Router Protocol on page 232](#)

Example: Configuring the PPPoE-Based Radio-to-Router Protocol

This example shows how to configure the PPPoE-based radio-to-router protocol.

- [Requirements on page 233](#)
- [Overview on page 233](#)
- [Configuration on page 233](#)
- [Verification on page 234](#)

Requirements

Before you begin:

1. Configure network interfaces. See “[Example: Creating an Ethernet Interface](#)” on [page 122](#).
2. Configure PPPoE interfaces. See “[Example: Configuring PPPoE Interfaces](#)” on [page 211](#).
3. Configure PPPoE encapsulation on an Ethernet interface. See “[Example: Configuring PPPoE Encapsulation on an Ethernet Interface](#)” on [page 217](#).
4. Configure PPPoE encapsulation on an ATM-over-ADSL interface. See “[Example: Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface](#)” on [page 219](#).
5. Configure CHAP authentication on a PPPoE interface. See “[Example: Configuring CHAP Authentication on a PPPoE Interface](#)” on [page 221](#).

Overview

In this example, you configure the ge-3/0/3 interface and set the bandwidth, resource, latency, and quality to **100**. You also set the threshold value to **10**, and then configure options on the logical interface.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following command, paste it into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the command into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
[edit]
set interfaces ge-3/0/3 unit 1 radio-router bandwidth 100 resource 100 latency 100 quality
  100 threshold 10
set interfaces pp0 unit 1 pppoe-options underlying-interface ge-3/0/3 server
set interfaces pp0 unit 1 family inet unnumbered-address lo0.0 destination 192.168.1.2
set interfaces pp0 unit 1 family inet6 address lo0.0 destination fec0:1:1::2
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see the [Junos OS CLI User Guide](#).

To configure the PPPoE-based radio-to-router protocol:

1. Enable the PPPoE-based radio-to-router protocol.

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

2. Set the interface speed for the virtual link.

```
[edit interfaces ge-3/0/3 unit 1 radio-router]
user@host# set bandwidth 100 resource 100 latency 100 quality 100
```

3. Set the calculated and current interface speeds, as a percentage.

```
[edit interfaces ge-3/0/3 unit 1 radio-router]
user@host# set threshold 10
```

4. Configure options on the logical interface.

```
[edit interfaces pp0 unit 1]
user@host# set pppoe-options underlying-interface ge-3/0/3
user@host# set pppoe-options server
user@host# set family inet unnumbered-address lo0.0 destination 192.168.1.2
user@host# set family inet6 address lo0.0 destination fec0:1:1::2
```

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

For brevity, this **show interfaces** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
user@host# show interfaces ge-3/0/3 {
  unit 1
  radio-router {
    bandwidth 100;
    resource 100;
    latency 100;
    quality 100;
    threshold 10;
  }
}
...
pp0 {
  unit 1 {
    pppoe-options {
      underlying-interface ge-3/0/3;
      server;
    }
  }
  family inet {
    unnumbered-address lo0.0 destination 192.168.1.2;
  }
  family inet6;
}
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Verifying the PPPoE-based Radio-to-Router Protocol

Purpose Verify the PPPoE-Based radio-to-router protocol.

Action From operational mode, enter the **show interfaces** command.

- Related Documentation**
- [Understanding the PPPoE-Based Radio-to-Router Protocol on page 230](#)

Credit Flow Control for PPPoE

To support the credit-based flow control extensions described in RFC4938, PPPoE peers can grant each other forwarding credits. The grantee is allowed to forward traffic to the peer only when it has a sufficient number of credits to do so. When credit-based forwarding is used on both sides of the session, the radio client can throttle traffic by limiting the number of credits it grants to the router.

The **interfaces** statement includes the **radio-router** attribute, which contains the parameters used for rate-based scheduling and OSPF link cost calculations. It also includes the **credit** attribute to indicate that credit-based packet scheduling is supported on the PPPoE interfaces that reference this underlying interface. Interfaces that set the **encapsulation** attribute support the PPPoE Active Discovery Grant (PADG) and PPPoE Active Discovery Credit (PADC) messages in the same way that the **radio-router** attribute provides active support for the PPPoE Active Discovery Quality (PADQ) message.

The **credit interval** parameter controls how frequently the router generates credit announcement messages. For PPPoE this corresponds to the interval between PADG credit announcements for each session.

- Related Documentation**
- [PPPoE-Based Radio-to-Router Protocols Overview on page 229](#)
 - [Understanding the PPPoE-Based Radio-to-Router Protocol on page 230](#)
 - [Configuring PPPoE-Based Radio-to-Router Protocols on page 232](#)

PPPoE Credit-Based Flow Control Configuration

This example shows a PPPoE credit-based flow control configuration.

```
[edit interfaces ge-0/0/1]
unit 0 {
  encapsulation ppp-over-ether;
  radio-router {
    credit {
      interval 10;
    }
    bandwidth 80;
    threshold 5;
  }
}
```

- Related Documentation**
- [Understanding the PPPoE-Based Radio-to-Router Protocol on page 230](#)
 - [Configuring PPPoE-Based Radio-to-Router Protocols on page 232](#)

Configuring R2CP Radio-to-Router Protocol

- [R2CP Radio-to-Router Protocol Overview on page 237](#)
- [Configuring the R2CP Radio-to-Router Protocol on page 238](#)

R2CP Radio-to-Router Protocol Overview

Supported Platforms [SRX Series](#)

The Network Centric Waveform (NCW) radio-specific radio-to-router control protocol (R2CP) is similar to the PPPoE radio-to-router protocol. Both of these protocols exchange dynamic metric changes in the network that the routers use to update the OSPF topologies.

In radio-router topologies, the router connects to the radio over a Gigabit Ethernet link and the radio transmits packets over the radio frequency (RF) link. The radio periodically sends metrics to the router, which uses RF link characteristics and other data to inform the router on the shaping and OSPF link capacity. The router uses this information to shape the data traffic and provide the OSPF link cost for its SPF calculations. The radio functions like a Layer 2 switch and can only identify remote radio-router pairs using the Layer 2 MAC addresses. With R2CP the router receives metrics for each neighboring router, identified by the MAC address of the remote router. The R2CP daemon translates the MAC addresses to link the local IPv6 address and sends the metrics for each neighbor to OSPF. Processing these metrics is similar to the handling of PPPoE PADQ metrics. Unlike PPPoE, which is a point-to-point link, these R2CP neighbors are treated as nodes in a broadcast LAN.

You must configure each neighbor node with a per unit scheduler for CoS. The scheduler context defines the attributes of Junos class-of-service. To define CoS for each radio, you can configure virtual channels to limit traffic. You need to configure virtual channels for as many remote radio-router pairs as there are in the network. You configure virtual channels on a logical interface. Each virtual channel can be configured to have a set of eight queues with a scheduler and an optional shaper. When the radio initiates the session with a peer radio-router pair, a new session is created with the remote MAC address of the router and the VLAN over which the traffic flows. Junos OS chooses from the list of free virtual channels and assigns the remote MAC and the eight CoS queues and the

scheduler to this remote MAC address. All traffic destined to this remote MAC address is subjected to the CoS that is defined in the virtual channel.

A virtual channel group is a collection of virtual channels. Each radio can have only one virtual channel group assigned uniquely. If you have more than one radio connected to the router, you must have one virtual channel group for each local radio-to-router pair. Although a virtual channel group is assigned to a logical interface, a virtual channel is not the same as a logical interface. The only features supported on a virtual channel are queuing, packet scheduling, and accounting. Rewrite rules and routing protocols apply to the entire logical interface.

All nodes in the R2CP network are in a broadcast LAN. The point-to-multipoint over LAN protocol supports advertising different bandwidth information for neighbors on a broadcast link. The network link is a point-to-multipoint link in the OSPFv3 link state database, which uses existing OSPF neighbor discovery to provide automatic discovery without configuration. It enables each node to advertise a different metric to every other node in the network to accurately represent the cost of communication. The **p2mp-over-lan** interface type under the OSPFv3 interface configuration enables you to configure the interface. OSPFv3 then uses LAN procedures for neighbor discovery and flooding, but represents the interface as point-to-multipoint in the link state database.

The interface type and router LSA are available under the following hierarchies:

[protocols ospf3 area *area-id* interface *interface-name*]

[routing-instances *routing-instances-name* protocols ospf3 area *area-id* interface *interface-name*]

For example:

```
protocols {
  ospf3 {
    area 0.0.0.0 {
      interface ge-0/0/2.0 {
        interface-type p2mp-over-lan;
      }
    }
  }
}
```

**Related
Documentation**

- [Configuring the R2CP Radio-to-Router Protocol on page 238](#)

Configuring the R2CP Radio-to-Router Protocol

Supported Platforms [SRX Series](#)

To configure the R2CP protocol:

1. Configure the interfaces.

The following example creates four logical interfaces on ge-0/0/2, using unit 52 for R2CP control messages and units 101-193 for data traffic. The **per-unit-scheduler** statement is required for R2CP.

```

interfaces {
  ge-0/0/2 {
    per-unit-scheduler;
    vlan-tagging;
    unit 52 {
      vlan-id 52;
      family inet {
        address 52.1.1.1/24;
      }
    }
    unit 101 {
      vlan-id 101;
      family inet {
        address 101.1.1.1/24;
      }
    }
    unit 102 {
      vlan-id 102;
      family inet {
        address 102.1.1.1/24;
      }
    }
    unit 103 {
      vlan-id 103;
      family inet {
        address 103.1.1.1/24;
      }
    }
  }
}

```

2. Configure the R2CP protocol.

The following example configures g2-0/0/2.52 as the interface for R2CP control messages, vg1 as the virtual-channel group, and ge-0/0/2.101-103 as data interfaces using the radio-interface statement.

```

protocols {
  r2cp {
    radio myRadio {
      interface ge-0/0/2.52;
      virtual-channel-group vg1;
      radio-interface ge-0/0/2.101;
      radio-interface ge-0/0/2.102;
      radio-interface ge-0/0/2.103;
    }
  }
}

```

3. Configure class of service.

The following example defines virtual-channels, their initial shaping-rates, and the virtual-channel-group to which they belong. It also makes the association between

radio-interface interfaces and virtual-channel-group. In the class of service configuration, the **vc-shared-scheduler** configuration statement is required for each interface configured as a radio interface in the R2CP protocol configuration.

```
class-of-service {
  virtual-channels {
    vc1;
    vc2;
    vc3;
    vc4;
  }
  virtual-channel-groups {
    vg1 {
      vc1 {
        scheduler-map sm;
        shaping-rate 15m;
        default;
      }
      vc2 {
        scheduler-map sm;
        shaping-rate 20m;
      }
      vc3 {
        scheduler-map sm;
        shaping-rate 20m;
      }
      vc4 {
        scheduler-map sm;
        shaping-rate 20m;
      }
    }
  }
  forwarding-classes {
    queue 0 DATA-queue;
  }
  interfaces {
    ge-0/0/2 {
      unit 101 {
        virtual-channel-group vg1;
        vc-shared-scheduler;
      }
      unit 102 {
        virtual-channel-group vg1;
        vc-shared-scheduler;
      }
      unit 103 {
        virtual-channel-group vg1;
        vc-shared-scheduler;
      }
    }
  }
  scheduler-maps {
    sm {
      forwarding-class DATA-queue scheduler sm-scheduler;
    }
  }
}
```

```
schedulers {  
  sm-scheduler {  
    transmit-rate percent 20;  
    buffer-size percent 20;  
    priority low;  
  }  
}
```

Related Documentation • [R2CP Radio-to-Router Protocol Overview on page 237](#)

PART 6

Configuration Statements and Operational Commands

- Configuration Statements on page 245
- Operational Commands on page 317

CHAPTER 18

Configuration Statements

- `accept-source-mac` on page 248
- `access-point-name` on page 249
- `apply-groups` on page 249
- `arp-resp` on page 250
- `authentication-method` (Interfaces) on page 250
- `bandwidth` (Interfaces) on page 251
- `bundle` (Interfaces) on page 251
- `cbr rate` on page 252
- `cellular-options` on page 252
- `classifiers` (CoS) on page 253
- `client-identifier` (Interfaces) on page 254
- `code-points` (CoS) on page 254
- `compression-device` (Interfaces) on page 255
- `credit` (Interfaces) on page 255
- `data-rate` on page 256
- `disable` (PoE) on page 256
- `dhcp` (Interfaces) on page 257
- `duration` (PoE) on page 258
- `encapsulation` (Interfaces) on page 259
- `family inet` (Interfaces) on page 260
- `family inet6` on page 263
- `flag` (Interfaces) on page 265
- `flexible-vlan-tagging` (Interfaces) on page 266
- `flow-control` (Interfaces) on page 266
- `flow-monitoring` (Services) on page 267
- `forwarding-classes` (CoS) on page 268
- `fpc` (Interfaces) on page 269
- `gratuitous-arp-reply` on page 270

- [gsm-options](#) on page 270
- [guard-band \(PoE\)](#) on page 271
- [hub-assist](#) on page 271
- [inline-jflow \(Forwarding Options\)](#) on page 272
- [interface \(PIC Bundle\)](#) on page 272
- [interface \(PoE\)](#) on page 273
- [interfaces \(CoS\)](#) on page 274
- [interval \(Interfaces\)](#) on page 275
- [interval \(PoE\)](#) on page 275
- [ipv4-template \(Services\)](#) on page 276
- [ipv6-template \(Services\)](#) on page 276
- [lACP \(Interfaces\)](#) on page 277
- [latency \(Interfaces\)](#) on page 277
- [lease-time](#) on page 278
- [line-rate \(Interfaces\)](#) on page 278
- [link-speed \(Interfaces\)](#) on page 279
- [loopback \(Interfaces\)](#) on page 279
- [loss-priority \(CoS Loss Priority\)](#) on page 280
- [loss-priority \(CoS Rewrite Rules\)](#) on page 281
- [loss-priority-maps \(CoS Interfaces\)](#) on page 282
- [loss-priority-maps \(CoS\)](#) on page 282
- [management \(PoE\)](#) on page 283
- [maximum-power \(PoE\)](#) on page 283
- [media-type \(Interfaces\)](#) on page 284
- [minimum-links \(Interfaces\)](#) on page 285
- [native-vlan-id \(Interfaces\)](#) on page 286
- [next-hop-tunnel](#) on page 286
- [no-dns-propagation](#) on page 287
- [option-refresh-rate \(Services\)](#) on page 287
- [pic-mode \(Chassis T1 Mode\)](#) on page 288
- [periodic \(Interfaces\)](#) on page 289
- [ppp-over-ether](#) on page 289
- [pppoe](#) on page 290
- [pppoe-options](#) on page 291
- [priority \(PoE\)](#) on page 292
- [profile \(Access\)](#) on page 293
- [profiles](#) on page 295

- promiscuous-mode (Interfaces) on page 296
- quality (Interfaces) on page 296
- r2cp on page 297
- radio-router (Interfaces) on page 298
- redundancy-group (Interfaces) on page 299
- redundant-ether-options on page 300
- redundant-parent (Interfaces Fast Ethernet) on page 301
- redundant-parent (Interfaces Gigabit Ethernet) on page 301
- resource (Interfaces) on page 302
- retransmission-attempt on page 302
- retransmission-interval (Interfaces) on page 303
- roaming-mode on page 303
- scheduler-map (CoS Virtual Channels) on page 304
- select-profile on page 304
- server-address on page 305
- shaping-rate (CoS Interfaces) on page 306
- simple-filter (Interfaces) on page 307
- sip-password on page 307
- sip-user-id on page 308
- source-address-filter (Interfaces) on page 309
- source-filtering (Interfaces) on page 310
- speed (Interfaces) on page 310
- telemetries (PoE) on page 311
- template-refresh-rate (Services) on page 311
- threshold (Interfaces) on page 312
- traceoptions (Interfaces) on page 312
- update-server on page 313
- vbr rate on page 313
- vdsl-profile on page 314
- vendor-id (Interfaces) on page 314
- vlan-tagging (Interfaces) on page 315
- web-authentication (Interfaces) on page 316

accept-source-mac

Supported Platforms SRX1500, SRX1500, SRX300, SRX320, SRX340, vSRX

Syntax

```
accept-source-mac {
    mac-address mac-address;
}
```

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

Release Information Statement introduced in Junos OS Release 11.4.

Description For Gigabit Ethernet (GE), Fast Ethernet (FE), or 10 Gigabit Ethernet (XE) interfaces, specify the MAC addresses from which the interface can receive packets. Ensure that you update the MAC address if the remote Ethernet card is replaced. Replacing the interface card changes the MAC address. If you do not update the MAC address, the interface cannot receive packets from the new card.



NOTE:

- Software-based MAC limiting is supported on SRX300, SRX320, and SRX340 devices. A maximum of 32 MAC addresses is supported per device.

Options *mac-address* —MAC address filter. You can specify the MAC address as six hexadecimal bytes in one of the following formats: *nn:nn:nn:nn:nn:nn* (for example, 00:11:22:33:44:55) or *nnnn:nnnn:nnnn* (for example, 0011.2233.4455). You can configure up to 32 source addresses. To specify more than one address, include multiple *mac-addresses* in the *source-address-filter* statement.

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

Related Documentation

- [Understanding Ethernet Interfaces on page 117](#)

access-point-name

Supported Platforms	SRX300, SRX320, vSRX
Syntax	access-point-name <i>apn</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> cellular-options gsm-options profiles <i>profile-name</i>]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Configure the access point name (APN) provided by the service provider for connection to a Global System for Mobile Communications (GSM) cellular network.
Options	<i>apn</i> —Access point name.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

apply-groups

Supported Platforms	MX Series, vSRX
Syntax	apply-groups;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> radio-router]
Release Information	Statement introduced in Junos OS Release 9.6. Statement modified in Junos OS Release 15.1.
Description	Apply the groups from which to inherit configuration data. If radio-router is set without any other attributes specified, the first four values become 100 and threshold stays at 10, and capacity, margin, and delay are deprecated. If radio-router is set, do not change the OSPF reference-bandwidth value because this generates an incorrect link cost.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring PPPoE-Based Radio-to-Router Protocols on page 232

arp-resp

Syntax	arp-resp (restricted unrestricted);
Hierarchy Level	[edit interfaces <i>interfaces-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in Junos OS Release 10.1.
Description	Configure Address Resolution Protocol (ARP) response on the interface.
Options	<ul style="list-style-type: none">• restricted—Enable restricted proxy ARP response on the interface. This is the default.• unrestricted—Enable unrestricted ARP response on the interface.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Interfaces Feature Guide for Security Devices</i>

authentication-method (Interfaces)

Supported Platforms	SRX300, SRX320, vSRX
Syntax	authentication-method (pap chap none);
Hierarchy Level	[edit interfaces <i>interface-name</i> cellular-options gsm-options profiles <i>profile-name</i>]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Specify the authentication method for connection to a Global System for Mobile Communications (GSM) cellular network.
Options	<ul style="list-style-type: none">• pap—Password Authentication Protocol.• chap—Challenge Handshake Authentication Protocol.• none—No authentication method is used.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

bandwidth (Interfaces)

Supported Platforms	SRX Series, vSRX
Syntax	bandwidth <i>bandwidth</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> radio-router]
Release Information	Statement introduced in Junos OS Release 10.1.
Description	This option controls the weight of the current (vs. maximum) data rate (value 0–100).
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • PPPoE-Based Radio-to-Router Protocols Overview on page 229

bundle (Interfaces)

Supported Platforms	SRX Series, vSRX
Syntax	bundle <i>bundle-name</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family mlppp]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Specify the logical interface name the link joins.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Interfaces on page 3

cbr rate

Supported Platforms	SRX Series
Syntax	cbr rate;
Hierarchy Level	[edit interfaces interface-name atm-options vpi vpi-identifier shaping]
Release Information	Command introduced in Release 9.5 of Junos OS.
Description	For ATM encapsulation only, define a constant bit rate bandwidth utilization in the traffic-shaping profile.
Options	<ul style="list-style-type: none"> • CBR Value—Constant bandwidth utilization (range: 33,000 through 1,199,920) • CDVT—Cell delay variation tolerance in microseconds (range: 1 through 9999)
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Junos OS Interfaces Configuration Guide for Security Devices</i>

cellular-options

Supported Platforms	SRX300, SRX320
Syntax	<pre>cellular-options { roaming-mode (home only automatic) gsm-options { select-profile profile-name; profiles { profile-name { sip-user-id simple-ip-user-id; sip-password simple-ip-password; access-point-name apn; authentication-method (pap chap none); } } } }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Configure options for connecting a 3G wireless modem interface to a cellular network.
Options	The remaining statements are explained separately. See CLI Explorer .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

classifiers (CoS)

Supported Platforms [SRX Series, vSRX](#)

Syntax

```
classifiers {
  (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence) classifier-name {
    forwarding-class forwarding-class-name {
      loss-priority (high | low | medium-high | medium-low) {
        code-point alias-or-bit-string ;
      }
    }
    import (default | user-defined);
  }
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced in Junos OS Release 9.2

Description Configure a user-defined behavior aggregate (BA) classifier.

- Options**
- *classifier-name*—User-defined name for the classifier.
 - `import (default | user-defined)`—Specify the template to use to map any code points not explicitly mapped in this configuration. For example, if the classifier is of type **dscp** and you specify **import default**, code points you do not map in your configuration will use the predefined DSCP default mapping; if you specify **import mymap**, for example, code points not mapped in the forwarding-class configuration would use the mappings in a user-defined classifier named **mymap**.
 - `forwarding-class class-name`—Specify the name of the forwarding class. You can use the default forwarding class names or define new ones.
 - `loss-priority level`—Specify a loss priority for this forwarding class: **high**, **low**, **medium-high**, **medium-low**.
 - `code-points (alias | bits)`—Specify a code-point alias or the code points that map to this forwarding class.

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

Related Documentation

- [Understanding Interfaces on page 3](#)

client-identifier (Interfaces)

Supported Platforms	SRX Series, vSRX
Syntax	client-identifier { (ascii <i>string</i> hexadecimal <i>string</i>); }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family-name</i> dhcp]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Specify an ASCII or hexadecimal identifier for the Dynamic Host Configuration Protocol (DHCP) client. The DHCP server identifies a client by a client-identifier value.
Options	<ul style="list-style-type: none"> • ascii <i>ascii</i>—Identifier consisting of ASCII characters. • hexadecimal <i>hexadecimal</i>—Identifier consisting of hexadecimal characters.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Interfaces on page 3

code-points (CoS)

Supported Platforms	SRX Series, vSRX
Syntax	code-points [<i>aliases</i>] [<i>bit-patterns</i>];
Hierarchy Level	[edit class-of-service classifiers (dscp) <i>classifier-name</i> forwarding-class <i>class-name</i> loss-priority <i>level</i>]
Release Information	Statement introduced in Junos OS Release 11.1.
Description	Configure one or more code-point aliases or bit sets to apply to a forwarding class.
Options	<p><i>aliases</i>—Name of the alias or aliases.</p> <p><i>bit-patterns</i>—Value of the code-point bits, in decimal form.</p>
Required Privilege Level	interfaces—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Interfaces on page 3 • Example: Configuring BA Classifiers on Transparent Mode Devices

compression-device (Interfaces)

Supported Platforms	SRX Series, vSRX
Syntax	compression-device <i>name</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit (Interfaces) <i>logical-unit-number</i>]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Specify the compression interface for voice services traffic.
Options	<i>name</i> —Name of the AC.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Interfaces on page 3

credit (Interfaces)

Supported Platforms	SRX Series, vSRX
Syntax	credit { interval <i>number</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> radio—router]
Release Information	Statement introduced in Junos OS Release 10.1.
Description	This parameter controls credit-based scheduling parameters and includes an interval option to set the grant rate interval to a value between 1–60 seconds.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Interfaces on page 3

data-rate

Supported Platforms	LN Series
Syntax	<code>data-rate weight;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> radio-router]
Release Information	Statement introduced in Release 10.2 of Junos OS .
Description	Configure the weight of the resource factor when calculating an effective data rate.
Options	<p>weight—Factor used to calculate data rate.</p> <p>Range: 0 through 100</p> <p>Default: 100</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring PPPoE-Based Radio-to-Router Protocols on page 232

disable (PoE)

Supported Platforms	SRX1500, SRX320, SRX340, SRX5400, SRX550M
Syntax	<code>disable;</code>
Hierarchy Level	<p>[edit poe interface (all <i>interface-name</i>)</p> <p>[edit poe interface (all <i>interface-name</i>) telemetries]</p>
Release Information	Statement introduced in Junos OS Release 9.5.
Description	<p>Disables the PoE capabilities of the port. If PoE capabilities are disabled for a port, the port operates as a standard network access port. If the disable statement is specified after the telemetries statement, logging of PoE power consumption for the port is disabled. To disable monitoring and retain the stored interval and duration values for possible future use, you can specify the disable sub statement in the sub stanza for telemetries. Similarly for retaining the port configuration but disabling the PoE feature on the port, disable can be used in sub stanza for interface.</p>
Default	The PoE capabilities are automatically enabled when a PoE interface is set. Specifying the telemetries statement enables monitoring of PoE per-port power consumption.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Disabling a PoE Interface on page 194

dhcp (Interfaces)

Supported Platforms SRX Series, vSRX

Syntax dhcp {
 client-identifier {
 (ascii *string* | hexadecimal *string*);
 }
 lease-time (*length* | infinite);
 retransmission-attempt *value*;
 retransmission-interval *seconds*;
 server-address *server-address*;
 update-server;
 vendor-id *vendor-id* ;
}

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

Release Information Statement introduced in Junos OS Release 9.2.

Description Configure the Dynamic Host Configuration Protocol (DHCP) client.

Options The statements are explained separately.

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

Related Documentation

- [Understanding Interfaces on page 3](#)

duration (PoE)

Supported Platforms	SRX1500, SRX320, SRX340, SRX5400, SRX550M
Syntax	duration <i>hours</i> ;
Hierarchy Level	[edit poe interface (all <i>interface-name</i>) telemetries]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Modifies the duration for which telemetry records are stored. If telemetry logging continues beyond the specified duration, the older records are discarded one by one as new records are collected.
Options	hours— Hours for which telemetry data should be retained. Range: 1 through 24 hours Default: 1 hour
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring PoE on All Interfaces on page 189

encapsulation (Interfaces)

Supported Platforms	SRX Series, vSRX
Syntax	encapsulation (ether-vpls-ppp ethernet-bridge ethernet-ccc ethernet-tcc ethernet-vpls extended-frame-relay-ccc extended-frame-relay-tcc extended-vlan-bridge extended-vlan-ccc extended-vlan-tcc extended-vlan-vpls frame-relay-port-ccc vlan-ccc vlan-vpls);
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Specify logical link layer encapsulation.
Options	<ul style="list-style-type: none"> • cisco-hdlc—For normal mode (when the device is using only one B-channel). Cisco-compatible High-Level Data Link Control is a group of protocols for transmitting data between network points • frame-relay—Configure a Frame Relay encapsulation when the physical interface has multiple logical units, and the units are either point to point or multipoint. • multilink-frame-relay-uni-nni—Link services interfaces functioning as FRF.16 bundles can use Multilink Frame Relay UNI NNI encapsulation. • ppp—For normal mode (when the device is using only one ISDN B-channel per call). Point-to-Point Protocol is for communication between two computers using a serial interface. • ppp-over-ether—This encapsulation is used for underlying interfaces of pp0 interfaces.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Physical Encapsulation on an Interface on page 199

family inet (Interfaces)

Supported Platforms [SRX Series, vSRX](#)

```
Syntax  inet {
        accounting {
            destination-class-usage;
            source-class-usage {
                input;
                output;
            }
        }
        address (source-address/prefix) {
            arp destination-address {
                (mac mac-address | multicast-mac multicast-mac-address);
                publish publish-address;
            }
            broadcast address;
            preferred;
            primary;
            vrrp-group group-id {
                (accept-data | no-accept-data);
                advertise-interval seconds;
                advertisements-threshold number;
                authentication-key key-value;
                authentication-type (md5 | simple);
                fast-interval milliseconds;
                inet6-advertise-interval milliseconds
                (preempt <hold-time seconds> | no-preempt );
                priority value;
                track {
                    interface interface-name {
                        bandwidth-threshold bandwidth;
                        priority-cost value;
                    }
                    priority-hold-time seconds;
                    route route-address{
                        routing-instance routing-instance;
                        priority-cost value;
                    }
                }
                virtual-address [address];
                virtual-link-local-address address;
                vrrp-inherit-from {
                    active-group value;
                    active-interface interface-name;
                }
            }
            web-authentication {
                http;
                https;
                redirect-to-https;
            }
        }
        dhcp {
```

```

client-identifier {
  (ascii string | hexadecimal string);
}
lease-time (length | infinite);
retransmission-attempt value;
retransmission-interval seconds;
server-address server-address;
update-server;
vendor-id vendor-id ;
}
dhcp-client {
  client-identifier {
    prefix {
      host-name;
      logical-system-name;
      routing-instance-name;
    }
    use-interface-description (device | logical);
    user-id (ascii string| hexadecimal string);
  }
  lease-time (length | infinite);
  retransmission-attempt value;
  retransmission-interval seconds;
  server-address server-address;
  update-server;
  vendor-id vendor-id ;
}
filter {
  group number;
  input filter-name;
  input-list [filter-name];
  output filter-name;
  output-list [filter-name];
}
mtu value;
no-neighbor-learn;
no-redirects;
policer {
  arp arp-name;
  input input-name;
  output output-name;
}
primary;
rpf-check {
  fail-filter filter-name;
  mode {
    loose;
  }
}
sampling {
  input;
  output;
  simple-filter;
}
targeted-broadcast {
  (forward-and-send-to-re | forward-only);
}

```

```

    }
    unnumbered-address {
      interface-name;
      preferred-source-address preferred-source-address;
    }
  }
}

```

Hierarchy Level [edit interfaces *interface* unit *unit*]

Release Information Statement introduced in a prior release of Junos OS.

Description Assign an IP address to a logical interface.

Options *ipaddress*—Specifies the IP address for the interface.



NOTE: You use family `inet` to assign an IPv4 address. You use family `inet6` to assign an IPv6 address. An interface can be configured with both an IPv4 and IPv6 address.

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

Related Documentation

- [Understanding Interfaces on page 3](#)

family inet6

Supported Platforms [SRX Series, vSRX](#)

```
Syntax inet6 {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            ouput;
        }
    }
}
address source-address/prefix {
    eui-64;
    ndp address {
        (mac mac-address | multicast-mac multicast-mac-address);
        publish;
    }
    preferred;
    primary;
    vrrp-inet6-group group_id {
        (accept-data | no-accept-data);
        advertisements-threshold number;
        authentication-key value;
        authentication-type (md5 | simple);
        fast-interval milliseconds;
        inet6-advertise-interval milliseconds;
        (preempt <hold-time seconds> | no-preempt );
        priority value;
        track {
            interface interface-name {
                bandwidth-threshold value;
                priority-cost value;
            }
            priority-hold-time seconds;
            route route-address{
                routing-instance routing-instance;
            }
        }
    }
    virtual-inet6-address [address];
    virtual-link-local-address address;
    vrrp-inherit-from {
        active-group value;
        active-interface interface-name;
    }
}
web-authentication {
    http;
    https;
    redirect-to-https;
}
}
(dad-disable | no-dad-disable);
dhcpv6-client {
    client-ia-type (ia-na | ia-pd);
```

```

client-identifier duid-type (duid-ll | duid-llt | vendor);
client-type (autoconfig | stateful);
rapid-commit;
req-option (dns-server | domain | fqdn | nis-domain | nis-server | ntp-server | sip-domain
  | sip-server |time-zone | vendor-spec);
retransmission-attempt number;
update-router-advertisement {
  interface interface-name;
}
update-server;
}
filter {
  group number;
  input filter-name;
  input-list [filter-name];
  output filter-name;
  output-list [filter-name];
}
mtu value;
nd6-stale-time seconds;
no-neighbor-learn;
policer {
  input input-name;
  output output-name;
}
rpf-check {
  fail-filter filter-name;
  mode {
    loose;
  }
}
sampling {
  input;
  output;
}
unnumbered-address {
  interface-name;
  preferred-source-address preferred-source-address;
}
}

```

Hierarchy Level [edit interfaces *interface* unit *unit*]

Release Information Statement supported in Junos 10.2 for SRX Series devices.

Description Assign an IP address to a logical interface.

Options *ipaddress*—Specifies the IP address for the interface.



NOTE: You use family inet6 to assign an IPv6 address. You use family inet to assign an IPv4 address. An interface can be configured with both an IPv4 and IPv6 address.

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

Related Documentation • [Understanding Interfaces on page 3](#)

flag (Interfaces)

Supported Platforms [SRX Series, vSRX](#)

Syntax flag

Hierarchy Level [edit interfaces interface-name traceoptions]

Release Information Statement introduced in Junos OS Release 10.1.

Description Define tracing operations for individual interfaces. To specify more than one tracing operation, include multiple flag statements.

- Options**
- **all**—Enable all interface trace flags.
 - **event** —Trace interface events.
 - **cache**—Enable interface flags for Web filtering cache maintained on the routing table.
 - **enhanced**—Enable interface flags for processing through Enhanced Web Filtering.
 - **ipc**—Trace interface IPC messages.
 - **media**—Trace interface media changes.
 - **critical**—Trace critical events.
 - **major**—Trace major events.



NOTE:

- MTU is limited to 1518 on this interface.
 - **Cache and enhanced** options are applicable only to Enhanced Web Filtering.
-

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

Related Documentation • [Understanding Interfaces on page 3](#)

flexible-vlan-tagging (Interfaces)

Supported Platforms	SRX1500, SRX300, SRX320, SRX340, vSRX
Syntax	flexible-vlan-tagging;
Hierarchy Level	[edit interfaces <i>interface</i>]
Release Information	Statement introduced in Junos OS Release 12.1X44-D10.
Description	Simultaneously supports transmission of 802.1Q VLAN single-tag and dual-tag frames on logical interfaces on the same Ethernet port.
Options	native-vlan-id —Configures a VLAN identifier for single-tag frames, dual-tag frames, or a mixture of single-tag and dual-tag frames.
Required Privilege Level	interface —To view this statement in the configuration. interface-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring VLAN Tagging on page 49

flow-control (Interfaces)

Supported Platforms	SRX Series, vSRX
Syntax	(flow-control no-flow-control);
Hierarchy Level	[edit interfaces <i>interface-name</i> fastether-options], [edit interfaces <i>interface-name</i> gigheter-options], [edit interfaces <i>interface-name</i> redundant-ether-options]
Release Information	Statement modified in Junos OS Release 9.2.
Description	For Fast Ethernet, Gigabit Ethernet, and redundant Ethernet interfaces only, explicitly enable flow control, which regulates the flow of packets from the device to the remote side of the connection. Enabling flow control is useful when the device is a Gigabit Ethernet switch.
Default	Flow control is the default behavior.
Required Privilege Level	interface —To view this statement in the configuration. interface-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Ethernet Interfaces on page 117

flow-monitoring (Services)

Supported Platforms [SRX Series, vSRX](#)

Syntax

```

flow-monitoring {
  version9 {
    template template-name {
      flow-active-timeout seconds;
      flow-inactive-timeout seconds;
      ipv4-template;
      ipv6-template;
      option-refresh-rate {
        packets packets;
        seconds seconds;
      }
      template-refresh-rate {
        packets packets;
        seconds seconds;
      }
    }
  }
}

```

Hierarchy Level [edit services]

Release Information Statement introduced in Junos OS Release 10.4.

Description Configure flow monitoring.

Options `version9`—Version 9 configuration.

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

Related Documentation

- [Understanding Interfaces on page 3](#)

forwarding-classes (CoS)

Supported Platforms SRX Series, vSRX

Syntax

```
forwarding-classes {
  class class-name {
    priority (high | low);
    queue-num number;
    spu-priority (high | low);
  }
  queue queue-number {
    class-name {
      priority (high | low);
    }
  }
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced in Junos OS Release 8.5. Statement updated in Junos OS Release 11.4. The **spu-priority** option introduced in Junos OS Release 11.4R2.

Description Configure forwarding classes and assign queue numbers.

Options

- *class-name*—Display the forwarding class name assigned to the internal queue number.



NOTE: This option is supported only on high-end SRX Series devices, including the SRX1500, SRX5400, SRX5600, and SRX5800.



NOTE: AppQoS forwarding classes must be different from those defined for interface-based rewriters.

- **policing-priority**—Layer 2 policing. One forwarding class can be configured as **premium** and others are configured as **normal**.
- **priority**—Fabric priority value:
 - **high**—Forwarding class's fabric queuing has high priority.
 - **low**—Forwarding class's fabric queuing has low priority.
- *queue-number*—Specify the internal queue number to which a forwarding class is assigned.
- **spu-priority**—Services Processing Unit (SPU) priority queue, either **high** or **low**.



NOTE: The `spu-priority` option is only supported on SRX1500 devices and SRX5000 line devices.

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

Related Documentation

- *Example: Configuring AppQoS*

fpc (Interfaces)

Supported Platforms vSRX

Syntax `fpc;`

Hierarchy Level [edit interfaces pic-set pic-set-name]

Release Information Command introduced in Junos OS Release 9.6.

Description Sets the PIC bundle and the FPC slot.

- Options**
- *apply-groups*—Inherit configuration data from these groups.
 - *apply-groups-except*—Do not inherit configuration data from these groups.

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

Related Documentation

- [Understanding Interfaces on page 3](#)

gratuitous-arp-reply

Syntax	(gratuitous-arp-reply no-gratuitous-arp-reply);
Hierarchy Level	[edit interfaces <i>interface-name</i>] [edit interfaces <i>interface-range</i> <i>interface-range-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For Ethernet interfaces, enable updating of the Address Resolution Protocol (ARP) cache for gratuitous ARPs.
Default	Updating of the ARP cache is disabled on all Ethernet interfaces.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Interfaces Feature Guide for Security Devices</i>

gsm-options

Supported Platforms	SRX300, SRX320
Syntax	<pre>gsm-options { select-profile <i>profile-name</i>; profiles { <i>profile-name</i> { sip-user-id <i>simple-ip-user-id</i>; sip-password <i>simple-ip-password</i>; access-point-name <i>apn</i>; authentication-method (pap chap none); } } }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> <i>cellular-options</i>]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Configure the 3G wireless modem interface to establish a data call with a Global System for Mobile Communications (GSM) cellular network.
Options	The remaining statements are explained separately. See CLI Explorer .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

guard-band (PoE)

Supported Platforms	SRX1500, SRX320, SRX340, SRX5400, SRX550M
Syntax	guard-band <i>watts</i> ;
Hierarchy Level	[edit poe]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Reserves the specified amount of power for the SRX Series device in case of a spike in PoE consumption.
Options	<i>watts</i> —Amount of power to be reserved for the SRX Series device in case of a spike in PoE consumption. Range: 0 through 19 W Default: 0 W
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Power over Ethernet on page 187

hub-assist

Syntax	hub-assist <i>weight</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> radio-router]
Release Information	Statement introduced in Junos OS Release 10.2.
Description	Configure the weight of the resource factor when calculating an effective interface bandwidth.
Options	<i>weight</i> —Factor used to calculate interface bandwidth. Range: 0 through 100 Default: 100
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring PPPoE-Based Radio-to-Router Protocols on page 232

inline-jflow (Forwarding Options)

Supported Platforms	SRX Series, vSRX
Syntax	<pre>inline-jflow { flow-export-rate <i>number</i>; source-address <i>ip-address</i>; }</pre>
Hierarchy Level	[edit forwarding-options sampling instance <i>instance-name</i> family inet output] [edit forwarding-options sampling instance <i>instance-name</i> family inet6 output]
Release Information	Statement introduced in Junos OS Release 10.4. Support for family inet6 added in Junos OS Release 12.1X45-D10.
Description	Specify Inline processing of sampled packets.
Options	<ul style="list-style-type: none">• flow-export-rate <i>value</i>—Flow export rate of monitored packets in kpps. The range is from 1 through 400.• source-address <i>address</i>—Address to use for generating monitored packets.
Required Privilege Level	services—To view this statement in the configuration. services-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Interfaces on page 3

interface (PIC Bundle)

Supported Platforms	vSRX
Syntax	<pre>interface <i>interface-name</i>;</pre>
Hierarchy Level	[edit interfaces pic-set pic-set-name]
Release Information	Command introduced in Junos OS Release 9.6.
Description	Sets the PIC bundle and the interface.
Options	<ul style="list-style-type: none">• apply-groups—Groups from which to inherit configuration data.• apply-groups-except—Do not inherit configuration data from these groups.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Interfaces on page 3

interface (PoE)

Supported Platforms SRX1500, SRX320, SRX340, SRX5400, SRX550M

Syntax interface (all | *interface-name*) {
 disable;
 maximum-power *watts*;
 priority (high | low);
 telemetries {
 disable;
 duration *hours*;
 interval *minutes*;
 }
 }

Hierarchy Level [edit poe]

Release Information Statement introduced in Junos OS Release 9.5.

Description Enable a PoE interface for a PoE port. The PoE interface must be enabled in order for the port to provide power to a connected powered device.

Default The PoE interface is enabled by default

- Options**
- **all**— Apply the configuration to all interfaces on the SRX Series device that have not been explicitly configured otherwise.
 - **interface-name**— Explicitly configure a specific interface.

The remaining statements are explained separately. See [CLI Explorer](#).

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

Related Documentation

- [Understanding Power over Ethernet on page 187](#)

interfaces (CoS)

```
Syntax  interfaces {
        interface-name {
            input-scheduler-map map-name ;
            input-shaping-rate rate ;
            scheduler-map map-name ;
            scheduler-map-chassis map-name ;
            shaping-rate rate ;
            unit logical-unit-number {
                adaptive-shaper adaptive-shaper-name ;
                classifiers {
                    (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence)
                    ( classifier-name | default);
                }
                forwarding-class class-name ;
                fragmentation-map map-name ;
                input-scheduler-map map-name ;
                input-shaping-rate (percent percentage | rate );
                input-traffic-control-profile profiler-name shared-instance instance-name ;
                loss-priority-maps {
                    default;
                    map-name ;
                }
                output-traffic-control-profile profile-name shared-instance instance-name ;
                rewrite-rules {
                    dscp ( rewrite-name | default);
                    dscp-ipv6 ( rewrite-name | default);
                    exp ( rewrite-name | default) protocol protocol-types ;
                    frame-relay-de ( rewrite-name | default);
                    inet-precedence ( rewrite-name | default);
                }
                scheduler-map map-name ;
                shaping-rate rate ;
                virtual-channel-group group-name ;
            }
        }
    }
```

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

Release Information Statement introduced in Junos OS Release 8.5.

Description Associate the class-of-service configuration elements with an interface.

Options interface *interface-name* unit *number*—The user-specified interface name and unit number.

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

Related Documentation

- *Class of Service Feature Guide for Security Devices*

interval (Interfaces)

Supported Platforms	LN Series
Syntax	interval <i>seconds</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> radio-router credit]
Release Information	Statement introduced in Release 10.1 of Junos OS.
Description	Configure the frequency that the router generates credit announcement messages.
Options	<p>seconds—Interval between PADG credit announcements for each session.</p> <p>Range: 0 through 60</p> <p>Default: 1</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring PPPoE-Based Radio-to-Router Protocols on page 232

interval (PoE)

Supported Platforms	SRX1500, SRX320, SRX340, SRX5400, SRX550M
Syntax	interval <i>minutes</i> ;
Hierarchy Level	[edit poe interface (all <i>interface-name</i>) telemetries]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Modifies the interval for logging telemetries if you are monitoring the per-port power consumption for PoE interfaces.
Options	<p>minutes—Interval at which data is logged.</p> <p>Range: 1 through 30 minutes</p> <p>Default: 5 minutes</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Understanding Interfaces on page 3

ipv4-template (Services)

Supported Platforms	SRX Series, vSRX
Syntax	ipv4-template;
Hierarchy Level	[edit services flow-monitoring version9 template <i>template-name</i>]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Specify that the flow monitoring version 9 template is used only for IPv4 records.
Required Privilege Level	services—To view this statement in the configuration. services-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Juniper Networks Devices Processing Overview• Understanding Interfaces on page 3

ipv6-template (Services)

Supported Platforms	SRX Series, vSRX
Syntax	ipv6-template;
Hierarchy Level	[edit services flow-monitoring version9 template <i>template-name</i>]
Release Information	Statement introduced in Junos OS Release 12.1X45-D10.
Description	Specify that the flow monitoring version 9 template is used only for IPv6 records.
Required Privilege Level	services—To view this statement in the configuration. services-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Juniper Networks Devices Processing Overview• Understanding Interfaces on page 3

lACP (Interfaces)

Supported Platforms	SRX Series
Syntax	<pre>lACP { active; passive; periodic; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> redundant-ether-options]
Release Information	Statement introduced in Junos OS Release 10.2.
Description	For redundant Ethernet interfaces in a chassis cluster only, configure Link Aggregation Control Protocol (LACP).
Options	<ul style="list-style-type: none"> • active—Initiate transmission of LACP packets. • passive—Respond to LACP packets. • periodic—Interval for periodic transmission of LACP packets. <p>Default: If you do not specify lACP as either active or passive, LACP remains off (the default).</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Understanding LACP on Standalone Devices on page 149

latency (Interfaces)

Supported Platforms	SRX Series , vSRX
Syntax	latency <i>number</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> radio—router]
Release Information	Statement introduced in Junos OS Release 10.1.
Description	This option controls the latency weight (value 0–100).
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • PPPoE-Based Radio-to-Router Protocols Overview on page 229

lease-time

Supported Platforms	EX Series, QFX Series, SRX Series, vSRX
Syntax	lease-time (<i>length</i> infinite);
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet dhcp]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Request a specific lease time for the IP address.
Default	If no lease time is requested by client, then the server sends the lease time. The default lease time on a Junos OS DHCP server is one day.
Options	seconds —Request a lease time of a specific duration. Range: 60 through 2147483647 seconds infinite —Request that the lease never expire.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Interfaces on page 3

line-rate (Interfaces)

Supported Platforms	SRX Series, vSRX
Syntax	line-rate
Hierarchy Level	[edit interfaces <i>interfaces name</i> shdsl-options]
Release Information	Command introduced in Junos OS Release 10.0.
Description	Specify a line rate for an G.SHDSL interface.
Options	<ul style="list-style-type: none">• auto— Automatically selects a line rate.• value — Select the values between 192 kbps and 22784 kbps for the speed of transmission of data on the G.SHDSL connection.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring the G.SHDSL Interface on SRX Series Devices</i>

link-speed (Interfaces)

Syntax	<code>link-speed <i>speed</i>;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> redundant-ether-options]</code>
Release Information	Statement modified in Release 9.0 of Junos OS.
Description	For redundant Ethernet interfaces in a chassis cluster only, set the required link speed.
Options	<i>speed</i> —For redundant Ethernet links, you can specify <i>speed</i> in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Junos OS Interfaces Configuration Guide for Security Devices

loopback (Interfaces)

Supported Platforms	SRX Series , vSRX
Syntax	<code>(loopback no-loopback);</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> redundant-ether-options]</code>
Release Information	Statement modified in Junos OS Release 9.2.
Description	For Fast Ethernet, Gigabit Ethernet, and redundant Ethernet interfaces, enable or disable loopback mode.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Interfaces on page 3

loss-priority (CoS Loss Priority)

Supported Platforms	SRX Series, vSRX
Syntax	loss-priority <i>level</i> code-points[<i>values</i>];
Hierarchy Level	[edit class-of-service loss-priority-maps frame-relay-de <i>map-name</i>]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Map CoS values to a loss priority.
Options	<i>level</i> can be one of the following: <ul style="list-style-type: none">• high—Packet has high loss priority.• medium-high—Packet has medium-high loss priority.• medium-low—Packet has medium-low loss priority.• low—Packet has low loss priority.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Interfaces on page 3

loss-priority (CoS Rewrite Rules)

Supported Platforms	SRX Series, vSRX
Syntax	loss-priority <i>level</i> ;
Hierarchy Level	[edit class-of-service rewrite-rules <i>type rewrite-name</i> forwarding-class <i>class-name</i>]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Specify a loss priority to which to apply a rewrite rule. The rewrite rule sets the code-point aliases and bit patterns for a specific forwarding class and packet loss priority (PLP). The inputs for the map are the forwarding class and the PLP. The output of the map is the code-point alias or bit pattern.
Options	<i>level</i> can be one of the following: <ul style="list-style-type: none">• high—The rewrite rule applies to packets with high loss priority.• low—The rewrite rule applies to packets with low loss priority.• medium-high—The rewrite rule applies to packets with medium-high loss priority.• medium-low—The rewrite rule applies to packets with medium-low loss priority.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Class of Service Feature Guide for Security Devices</i>

loss-priority-maps (CoS Interfaces)

Supported Platforms	SRX Series, vSRX
Syntax	<pre>loss-priority-maps { frame-relay-de (<i>map-name</i> default); }</pre>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Assign the loss priority map to a logical interface.
Options	<ul style="list-style-type: none"> default—Apply default loss priority map. The default map contains the following: <pre>loss-priority low code-point 0; loss-priority high code-point 1;</pre> map-name—Name of loss priority map to be applied.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Understanding Interfaces on page 3

loss-priority-maps (CoS)

Supported Platforms	SRX Series, vSRX
Syntax	<pre>loss-priority-maps { frame-relay-de <i>loss-priority-map-name</i> { loss-priority (high low medium-high medium-low) { code-points [<i>bit-string</i>]; } } }</pre>
Hierarchy Level	[edit class-of-service]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Map the loss priority of incoming packets based on CoS values.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Understanding Interfaces on page 3

management (PoE)

Supported Platforms	SRX1500, SRX320, SRX340, SRX5400, SRX550M
Syntax	management (class static);
Hierarchy Level	[edit poe]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Designates how the SRX Series device allocates power to the PoE ports.
Default	static
Options	<ul style="list-style-type: none"> • static—When a powered device is connected to a PoE port, the power allocated to it is equal to the maximum power configured for the port. • class—When a powered device is connected to a PoE port, the power allocated to it is equal to the maximum power for the class as defined by the IEEE 802.3 AF standard.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring PoE on All Interfaces on page 189

maximum-power (PoE)

Supported Platforms	SRX1500, SRX320, SRX340, SRX5400, SRX550M
Syntax	maximum-power watts;
Hierarchy Level	[edit poe interface (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Maximum amount of power that can be supplied to the port.
Default	15.4 W
Options	<p>Watts—The maximum number of watts that can be supplied to the port.</p> <p>Range —0 through 15.4</p> <p>Default—15.4 W</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring PoE on All Interfaces on page 189

media-type (Interfaces)

Supported Platforms [SRX1500, SRX550M](#)

Syntax media-type

Hierarchy Level [edit interfaces *interface-name* media-type]

Release Information Command introduced in Junos OS Release 10.2.

Description Configure the operating modes for the 2-Port 10 Gigabit Ethernet XPIM.

- Options**
- copper
 - fiber

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

Related Documentation

- [Understanding Interfaces on page 3](#)

minimum-links (Interfaces)

Syntax	<code>minimum-links <i>number</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> redundant-ether-options]
Release Information	Statement added in Release 10.1 of Junos OS.
Description	<p>For redundant Ethernet interfaces configured as 802.3ad redundant Ethernet interface link aggregation groups (LAGs) in a chassis cluster only, set the required minimum number of physical child links on the primary node that must be working to prevent the interface from being down. Interfaces configured as redundant Ethernet interface LAGs typically have between 4 and 16 physical interfaces, but only half, those on the primary node, are relevant to the minimum-links setting.</p> <p>If the number of operating interfaces on the primary node falls below the configured value, it will cause the interface to be down even if some of the interfaces are still working.</p>
Options	<p><i>number</i>—For redundant Ethernet interface link aggregation group links, specify the number of physical child links on the primary node in the redundant Ethernet interface that must be working. The default minimum-links value is 1. The maximum value is half of the total number of physical child interfaces bound to the redundant Ethernet interface being configured or 8, whichever is smaller.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Junos OS Interfaces Configuration Guide for Security Devices</i>

native-vlan-id (Interfaces)

Supported Platforms	SRX Series, vSRX
Syntax	native-vlan-id <i>vlan-id</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Configure VLAN identifier for untagged packets received on the physical interface of a trunk mode interface.
Options	<i>vlan-id</i> —Configure a VLAN identifier for untagged packets. Enter a number from 0 through 4094.
Required Privilege Level	<i>interface</i> —To view this statement in the configuration. <i>interface-control</i> —To add this statement to the configuration
Related Documentation	<ul style="list-style-type: none"> • Understanding Interfaces on page 3

next-hop-tunnel

Supported Platforms	SRX Series, vSRX
Syntax	next-hop-tunnel <i>gateway-address</i> ipsec-vpn <i>vpn-name</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family-name</i>]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	For the secure tunnel (<i>st</i>) interface, create entries in the Next-Hop Tunnel Binding (NHTB) table, which is used to map the next-hop gateway IP address to a particular IP Security (IPsec) Virtual Private Network (VPN) tunnel. NHTB allows the binding of multiple IPsec VPN tunnels to a single IPsec tunnel interface.
Options	<ul style="list-style-type: none"> • <i>gateway-address</i>—Next-hop gateway IP address. • <i>ipsec-vpn vpn-name</i> —VPN to which the next-hop gateway IP address is mapped.
Required Privilege Level	<i>interface</i> —To view this statement in the configuration. <i>interface-control</i> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Interfaces on page 3

no-dns-propagation

Supported Platforms	SRX Series
Syntax	no-dns-propagation;
Hierarchy Level	[edit interface <i>interface-name</i> unit <i>unit-number</i> family <i>inet</i> <i>inet6</i> dhcp-client]
Release Information	Statement introduced in Junos OS Release 12.1X47-D35.
Description	Disable the propagation of DNS information to the kernel.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Interfaces on page 3

option-refresh-rate (Services)

Supported Platforms	SRX Series , vSRX
Syntax	option-refresh-rate
Hierarchy Level	[edit services flow-monitoring version9 template <i>template-name</i>]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Specify the option refresh rate.
Options	<ul style="list-style-type: none"> • packets—Specify the number of packets. The range is from 1 through 480,000. • seconds—Specify the number of seconds. The range is from 10 through 600.
Required Privilege Level	services—To view this statement in the configuration. services-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Flow Aggregation to Use Version 9 Flow Templates on page 25

pic-mode (Chassis T1 Mode)

Supported Platforms	SRX1500
Syntax	pic-mode (clear-channel);
Hierarchy Level	[edit chassis fpc <i>slot-number</i> pic <i>pic-number</i> ethernet]
Release Information	Statement added in Junos OS Release 10.2.
Description	Configure normal T1 mode or channelized T1 mode.
Options	<ul style="list-style-type: none"> • clear-channel—(default) Normal T1 mode. • ct1—Channelized T1 mode.



NOTE: When chassis clustering is enabled, it is necessary to indicate in the command which node is being configured. In such circumstances, the **edit chassis fpc** command becomes **edit chassis node *node-id* fpc**.

Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Interfaces on page 3

periodic (Interfaces)

Supported Platforms	SRX Series, vSRX
Syntax	periodic (fast slow);
Hierarchy Level	[edit interfaces <i>interface-name</i> redundant-ether-options lacp]
Release Information	Statement introduced in Junos OS Release 10.2.
Description	For redundant Ethernet interfaces in a chassis cluster only, configure the interval at which the interfaces on the remote side of the link transmit link aggregation control protocol data units (PDUs) by configuring the periodic statement on the interfaces on the local side. It is the configuration on the local side that specifies the behavior of the remote side. That is, the remote side transmits link aggregation control PDUs at the specified interval.
Options	<ul style="list-style-type: none"> fast—Transmit link aggregation control PDUs every second. slow—Transmit link aggregation control PDUs every 30 seconds.
	Default: fast
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Understanding Ethernet Interfaces on page 117

ppp-over-ether

Supported Platforms	SRX1500, SRX300, SRX320, SRX340, SRX550M
Syntax	ppp-over-ether;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> encapsulation]
Release Information	Statement introduced before Junos OS Release 11.2. This encapsulation is supported for Redundant Ethernet interface in Junos OS Release 11.2.
Description	This encapsulation is used for underlying interfaces of pp0 interfaces. This encapsulation is supported on Fast Ethernet interface, Gigabit Ethernet interface, and Redundant Ethernet interface. When Redundant Ethernet interface is used as underlying interface, an existing pppoe session can be continued in case of failover.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Understanding Ethernet Interfaces on page 117

pppoe

Supported Platforms [SRX Series](#)

Syntax

```
pppoe {  
    command binary-file-path;  
    disable;  
    failover (alternate-media | other-routing-engine);  
}
```

Hierarchy Level [edit system processes]

Release Information Statement introduced in Junos OS Release 9.2.

Description Enable users to connect to a network of hosts over a bridge or access concentrator.

- Options**
- **command *binary-file-path***—Path to the binary process.
 - **disable**—Disable the Point-to-Point Protocol over Ethernet process.
 - **failover**—Configure the device to reboot if the software process fails four times within 30 seconds, and specify the software to use during the reboot.
 - **alternate-media**—Configure the device to switch to backup media that contains a version of the system if a software process fails repeatedly.
 - **other-routing-engine**—Instruct the secondary Routing Engine to take mastership if a software process fails. If this statement is configured for a process, and that process fails four times within 30 seconds, then the device reboots from the secondary Routing Engine.

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

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

Related Documentation

- [Understanding Ethernet Interfaces on page 117](#)

pppoe-options

Supported Platforms	SRX1500, SRX300, SRX320, SRX340, SRX550M
Syntax	<pre>pppoe-options { access-concentrator <i>name</i> ; auto-reconnect <i>seconds</i>; (client server); ignore-eol-tag; service-name <i>name</i>; underlying-interface <i>interface-name</i>; }</pre>
Hierarchy Level	[edit interfaces pp0 unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces pp0 unit <i>logical-unit-number</i>]
Release Information	Statement modified in Junos OS Release 12.3X48 to include ignore-eol-tag statement.
Description	Configure PPP over Ethernet-specific interface properties.
Options	<p>access-concentrator <i>name</i>—(SRX Series devices with Point-to-Point Protocol over Ethernet (PPPoE) interfaces) Configure the name of the access concentrator. If you configure a specific access concentrator name on the client and the same access concentrator name server is available, then a PPPoE session is established. If there is a mismatch between the access concentrator names of the client and the server, the PPPoE session gets closed.</p> <p>auto-reconnect <i>seconds</i>—Configure the amount of time to wait before reconnecting after a session has terminated.</p> <p>client —Configure the device to operate in the PPPoE client mode.</p> <p>idle-timeout <i>seconds</i>—Configure the maximum time that a session can be idle.</p> <p>ignore-eol-tag—Disable the End-of-List tag to process the tags after the End-of-List tag in a PPPoE Active Discovery Offer (PADO) packet.</p> <p>service-name <i>name</i>—Configure the service to be requested from the PPP over Ethernet server; that is, the access concentrator. For example, you can use this statement to indicate an Internet service provider (ISP) name or a class of service.</p> <p>server—Configure the device to operate in the PPPoE server mode.</p> <p>underlying-interface <i>interface-name</i>—Configure the interface on which PPP over Ethernet is running.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring PPPoE Interfaces on page 211

priority (PoE)

Supported Platforms	SRX1500, SRX320, SRX340, SRX5400, SRX550M
Syntax	priority (high low);
Hierarchy Level	[edit poe interface (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Sets the priority of individual ports. When it is not possible to maintain power to all connected ports, lower-priority ports are powered off before higher priority ports. When a new device is connected on a higher-priority port, a lower-priority port will be powered off automatically if available power is insufficient to power on the higher-priority port. Note that for ports with the same priority configuration, ports on the left are given higher priority than the ports on the right.
Default	low
Options	value—high or low: <ul style="list-style-type: none">• high—Specify that this port is to be treated as high priority in terms of power allocation• low—Specify that this port is to be treated as low priority in terms of power allocation.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring PoE on All Interfaces on page 189

profile (Access)

Supported Platforms [SRX Series, vSRX](#)

```
Syntax profile profile-name {
  accounting {
    accounting-stop-on-access-deny;
    accounting-stop-on-failure;
    coa-immediate-update;
    duplication;
    immediate-update;
    order [accounting-method];
    statistics (time | volume-time);
    update-interval minutes;
  }
  accounting-order [accounting-method];
  address-assignment pool pool-name;
  authentication-order [ldap | none | password | securid];
  authorization-order [jsrc];
  client client-name {
    chap-secret chap-secret;
    client-group [group-names ];
    firewall-user {
      password password;
    }
    no-rfc2486;
    pap-password pap-password;
    x-auth ip-address;
  }
  client-name-filter {
    count number;
    domain-name domain-name;
    separator special-character;
  }
  ldap-options {
    assemble {
      common-name common-name;
    }
    base-distinguished-name base-distinguished-name;
    revert-interval seconds;
    search {
      admin-search {
        distinguished-name distinguished-name;
        password password;
      }
      search-filter search-filter-name;
    }
  }
  ldap-server server-address {
    port port-number;
    retry attempts;
    routing-instance routing-instance-name;
    source-address source-address;
    timeout seconds;
  }
}
```

```

provisioning-order (gx-plus | jsrc);
service {
  accounting-order {
    activation-protocol;
    radius;
  }
}
session-options {
  client-group [group-name];
  client-idle-timeout minutes;
  client-session-timeout minutes;
}
}

```

Hierarchy Level [edit access]

Release Information Statement introduced in Junos OS Release 10.4.

Description Create a profile containing a set of attributes that define device management access.

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

- Related Documentation**
- [Understanding Interfaces on page 3](#)
 - *Understanding User Authentication for Security Devices*
 - *Layer 2 Transparent and Switching Overview*

profiles

Supported Platforms SRX300, SRX320

Syntax

```
profiles {
  profile-name {
    sip-user-id simple-ip-user-id;
    sip-password simple-ip-password;
    access-point-name apn;
    authentication-method (pap | chap | none);
  }
}
```

Hierarchy Level [edit interfaces *interface-name* cellular-options gsm-options]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure a profile to establish a data call with a Global System for Mobile Communications (GSM) cellular network. You can configure up to 16 profiles.

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

The remaining statements are explained separately. See [CLI Explorer](#).

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

Related Documentation

- [Understanding Ethernet Interfaces on page 117](#)

promiscuous-mode (Interfaces)

Supported Platforms	SRX Series, vSRX
Syntax	promiscuous-mode;
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 10.1.
Description	<p>Enable promiscuous mode on Layer 3 Ethernet interfaces. When promiscuous mode is enabled on an interface, all packets received on the interface are sent to the central point or Services Processing Unit regardless of the destination MAC address of the packet.</p> <p>You can also enable promiscuous mode on chassis cluster redundant Ethernet interfaces and on aggregated Ethernet interfaces. If you enable promiscuous mode on a redundant Ethernet interface, promiscuous mode is then enabled on any child physical interfaces. If you enable promiscuous mode on an aggregated Ethernet interface, promiscuous mode is then enabled on all member interfaces.</p>
Required Privilege Level	interface —To view this statement in the configuration. interface-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Enabling and Disabling Promiscuous Mode on Ethernet Interfaces (CLI Procedure) on page 127

quality (Interfaces)

Supported Platforms	SRX Series, vSRX
Syntax	quality <value>;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> radio—router]
Release Information	Statement introduced in Junos OS Release 10.1.
Description	This option controls relative link quality weight (value 0–100).
Required Privilege Level	interface —To view this statement in the configuration. interface-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• PPPoE-Based Radio-to-Router Protocols Overview on page 229

r2cp

Supported Platforms	SRX1500, SRX300, SRX320, SRX340, SRX550M
Syntax	<pre>r2cp { command <i>binary-file-path</i>; disable; }</pre>
Hierarchy Level	[edit system processes]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Specify the Radio-to-Router Control Protocol (R2CP) used to exchange dynamic metric changes in the network that routers use to update the OSPF topologies.
Options	<ul style="list-style-type: none">• command <i>binary-file-path</i>—Path to the binary process.• disable—Disable the Radio-to-Router Control Protocol process.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• PPPoE-Based Radio-to-Router Protocols Overview on page 229

radio-router (Interfaces)

Supported Platforms [SRX Series](#)

Syntax

```
radio-router {  
  bandwidth number;  
  credit {  
    interval number;  
  }  
  data-rate number;  
  latency number;  
  quality number;  
  resource number;  
  threshold number;  
}
```

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

Release Information Statement introduced in Junos OS Release 10.1.

Description Point-to-Point Protocol over Ethernet (PPPoE)-based radio-to-router protocols include messages that define how an external system will provide the device with timely information about the quality of a link's connection. They also include a flow control mechanism to indicate how much data the device can forward. The device can then use the information provided in the PPPoE messages to dynamically adjust the interface speed of PPP links.

Options The remaining statements are explained separately. See [CLI Explorer](#).

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

Related Documentation

- [PPPoE-Based Radio-to-Router Protocols Overview on page 229](#)

redundancy-group (Interfaces)

Syntax	<code>redundancy-group <i>number</i> ;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> redundant-ether-options]
Release Information	Statement introduced in Junos OS Release 9.0.
Description	Specify the redundancy group that a redundant Ethernet interface belongs to.
Options	<i>number</i> —Number of the redundancy group that the redundant interface belongs to. Failover properties of the interface are inherited from the redundancy group. Range: 1 through 255
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Interfaces Feature Guide for Security Devices

redundant-ether-options

Supported Platforms [SRX Series, vSRX](#)

Syntax

```
redundant-ether-options {
  (flow-control | no-flow-control);
  lACP {
    (active | passive);
    periodic (fast | slow);
  }
  link-speed speed;
  (loopback | no-loopback);
  minimum-links number;
  redundancy-group number;
  source-address-filter mac-address;
  (source-filtering | no-source-filtering);
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Configure Ethernet redundancy options for a chassis cluster.

Options The remaining statements are explained separately. See [CLI Explorer](#).

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

Related Documentation

- *Example: Enabling Eight Queue Class of Service on Redundant Ethernet Interfaces*
- *Example: Configuring Chassis Cluster Redundant Ethernet Interfaces for IPv4 and IPv6 Addresses*

redundant-parent (Interfaces Fast Ethernet)

Supported Platforms	SRX Series, vSRX
Syntax	<code>redundant-parent interface-name ;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> fastether-options]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Configure Fast Ethernet-specific interface properties for Ethernet redundancy in a chassis cluster.
Options	<i>interface</i> —Parent redundant interface of the Fast Ethernet interface.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Ethernet Interfaces on page 117

redundant-parent (Interfaces Gigabit Ethernet)

Syntax	<code>redundant-parent interface-name ;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> ggether-options]
Release Information	Statement introduced in Release 9.0 of Junos OS.
Description	Configure Gigabit Ethernet-specific interface properties for Ethernet redundancy in a chassis cluster.
Options	<i>interface</i> —Parent redundant interface of the Gigabit Ethernet interface.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Junos OS Interfaces Configuration Guide for Security Devices

resource (Interfaces)

Supported Platforms	SRX Series, vSRX
Syntax	resource <i>number</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> radio—router]
Release Information	Statement introduced in Junos OS Release 10.1.
Description	This option controls the resource weight (value 1–100).
Required Privilege Level	interface —To view this statement in the configuration. interface-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • PPPoE-Based Radio-to-Router Protocols Overview on page 229

retransmission-attempt

Supported Platforms	EX Series, J Series, QFX Series, SRX Series
Syntax	retransmission-attempt <i>number</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet dhcp]
Release Information	Statement introduced in Junos OS Release 8.5 for J Series devices. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 9.2 for SRX Series devices. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the number of times the device retransmits a Dynamic Host Control Protocol (DHCP) packet if a DHCP server fails to respond. After the specified number of attempts, no further attempts at reaching a server are made.
Options	number —Number of retransmit attempts. Range: 0 through 6 Default: 4
Required Privilege Level	interface —To view this statement in the configuration. interface-control —To add this statement to the configuration.

retransmission-interval (Interfaces)

Syntax	retransmission-interval <i>seconds</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family-name</i> dhcp]
Release Information	Statement introduced in Release 8.5 of Junos OS.
Description	Specify the time between successive retransmission attempts.
Options	<p><i>seconds</i> —Number of seconds between successive retransmission.</p> <p>Range: 4 through 64 seconds</p> <p>Default: 4 seconds</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Junos OS Initial Configuration Guide for Security Devices</i>

roaming-mode

Supported Platforms	SRX320
Syntax	roaming-mode (home-only automatic)
Hierarchy Level	[edit interfaces <i>interface-name</i> cellular-options]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Specify whether the 3G wireless modem interface can access networks other than the home network.
Options	<ul style="list-style-type: none"> • home-only—No roaming is allowed. • automatic—Allows access to networks other than the home network. This is the default.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Understanding Ethernet Interfaces on page 117

scheduler-map (CoS Virtual Channels)

Supported Platforms	SRX Series, vSRX
Syntax	<code>scheduler-map map-name;</code>
Hierarchy Level	[edit class-of-service virtual-channel-groups <i>group-name</i> <i>virtual-channel-name</i>]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Apply a scheduler map to this virtual channel.
Options	<i>map-name</i> —Name of the scheduler map. The remaining statements are explained separately. See CLI Explorer .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>default (CoS)</i>• <i>shaping-rate (CoS Virtual Channels)</i>• <i>virtual-channel-group (CoS Interfaces)</i>• <i>virtual-channel-groups</i>• <i>virtual-channels</i>

select-profile

Supported Platforms	SRX Series, vSRX
Syntax	<code>select-profile profile-name</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> cellular-options gsm-options]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Select the active profile to establish a data call with a Global System for Mobile Communications (GSM) cellular network.
Options	<i>profile-name</i> —Name of a configured profile that is to be used to establish a data call.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Ethernet Interfaces on page 117

server-address

Supported Platforms	EX Series, QFX Series, SRX Series
Syntax	server-address <i>ip-address</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet dhcp]
Release Information	Statement introduced in Junos OS Release 8.5 for J Series devices. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 9.2 for SRX Series devices. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the address of the DHCP server that the client should accept DHCP offers from. If this option is included in the DHCP configuration, the client accepts offers only from this server and ignores all other offers.
Default	The client accepts the first offer it receives from any DHCP server.
Options	<i>ip-address</i> —DHCP server address.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

shaping-rate (CoS Interfaces)

Supported Platforms	SRX Series, vSRX
Syntax	shaping-rate <i>rate</i> ;
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i>], [edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	<p>For logical interfaces on which you configure packet scheduling, configure traffic shaping by specifying the amount of bandwidth to be allocated to the logical interface.</p> <p>Logical and physical interface traffic shaping is mutually exclusive. This means you can include the shaping-rate statement at the [edit class-of-service interfaces <i>interface interface-name</i>] hierarchy level or the [edit class-of-service interfaces <i>interface interface-name</i> unit <i>logical-unit-number</i>] hierarchy level, but not both.</p> <p>Alternatively, you can configure a shaping rate for a logical interface and oversubscribe the physical interface by including the shaping-rate statement at the [edit class-of-service traffic-control-profiles] hierarchy level. With this configuration approach, you can independently control the delay-buffer rate.</p>
Default	If you do not include this statement at the [edit class-of-service interfaces <i>interface interface-name</i> unit <i>logical-unit-number</i>] hierarchy level, the default logical interface bandwidth is the average of unused bandwidth for the number of logical interfaces that require default bandwidth treatment. If you do not include this statement at the [edit class-of-service interfaces <i>interface interface-name</i>] hierarchy level, the default physical interface bandwidth is the average of unused bandwidth for the number of physical interfaces that require default bandwidth treatment.
Options	<p>rate—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).</p> <p>Range: For logical interfaces, 1000 through 32,000,000,000 bps.</p> <p>For physical interfaces, 1000 through 160,000,000,000 bps.</p>
Required Privilege Level	<p><i>interface</i>—To view this statement in the configuration.</p> <p><i>interface-control</i>—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Class of Service Feature Guide for Security Devices

simple-filter (Interfaces)

Supported Platforms	SRX Series, vSRX
Syntax	simple-filter;
Hierarchy Level	[edit interfaces <i>interfaces-name</i> unit <i>logical-unit-number</i> family <i>family-name</i>]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Apply a simple filter to an interface. You can apply simple filters on ingress interfaces only.
Options	input <i>filter-name</i> : Name of one filter to evaluate when packets are received on the interface.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Ethernet Interfaces on page 117

sip-password

Supported Platforms	SRX300, SRX320
Syntax	sip-password <i>simple-ip-password</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> cellular-options gsm-options profiles <i>profile-name</i>]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Configure the password provided by the service provider for connection to a Global System for Mobile Communications (GSM) cellular network.
Options	<i>simple-ip-password</i> —Password.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Ethernet Interfaces on page 117

sip-user-id

Supported Platforms	SRX300, SRX320
Syntax	<code>sip-user-id simple-ip-user-id;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> cellular-options gsm-options profiles <i>profile-name</i>]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Configure the username provided by the service provider for connection to a Global System for Mobile Communications (GSM) cellular network.
Options	<i>simple-ip-user-id</i> —Username.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

source-address-filter (Interfaces)

Supported Platforms SRX1500, SRX300, SRX320, SRX340, SRX550M, vSRX

Syntax source-address-filter *mac-address* ;

Hierarchy Level [edit interfaces *interface-name* redundant-ether-options]

Release Information Statement modified in Junos OS Release 9.2.

Description For redundant Ethernet interfaces, specify the MAC addresses from which the interface can receive packets. For this statement to have any effect, you must include the **source-filtering** statement in the configuration to enable source address filtering.

Be sure to update the MAC address if the remote Ethernet card is replaced. Replacing the interface card changes the MAC address. Otherwise, the interface cannot receive packets from the new card.



NOTE:

- Software based MAC limiting is supported on SRX300, SRX320, and SRX340 devices.

A maximum of 32 devices are supported per device.

Options *mac-address* —MAC address filter. You can specify the MAC address as six hexadecimal bytes in one of the following formats: *nn:nn:nn:nn:nn:nn* (for example, 00:11:22:33:44:55) or *nxxx:xxxx:xxxx* (for example, 0011.2233.4455). You can configure up to 64 source addresses. To specify more than one address, include multiple *mac-address* options in the **source-address-filter** statement.

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

Related Documentation

- [Understanding Ethernet Interfaces on page 117](#)

source-filtering (Interfaces)

Supported Platforms	SRX1500, SRX300, SRX320, SRX340, SRX550M, vSRX
Syntax	(source-filtering no-source-filtering);
Hierarchy Level	[edit interfaces <i>interface-name</i> redundant-ether-options]
Release Information	Statement modified in Junos OS Release 9.2.
Description	<p>For redundant Ethernet interfaces, enable the filtering of MAC source addresses, which blocks all incoming packets to that interface. To allow the interface to receive packets from specific MAC addresses, include the source-address-filter statement.</p> <p>If the remote Ethernet card is changed, the interface cannot receive packets from the new card because it has a different MAC address.</p> <p>By default, source address filtering is disabled.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Understanding Ethernet Interfaces on page 117

speed (Interfaces)

Supported Platforms	SRX1500, SRX550M
Syntax	speed (100m 10m 1g);
Hierarchy Level	[edit interfaces <i>interface-name</i> speed]
Release Information	Command introduced in Junos OS Release 10.2.
Description	Configure the operating speed for the 2-Port 10 Gigabit Ethernet XPIM.
Options	<ul style="list-style-type: none"> • 100m — Link speed of 100 Mbps • 10g — Link speed of 10 Gbps • 10m — Link speed of 10 Mbps • 1g — Link speed of 1 Gbps
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Understanding Ethernet Interfaces on page 117 • Example: Configuring the 2-Port 10-Gigabit Ethernet XPIM Interface on page 176

telemetries (PoE)

Supported Platforms	SRX1500, SRX320, SRX340, SRX5400, SRX550M
Syntax	<pre>telemetries { disable; duration <i>hours</i>; interval <i>minutes</i>; }</pre>
Hierarchy Level	[edit poe interface (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Allow logging of per-port PoE power consumption. The telemetries section must be explicitly specified to enable logging. If left unspecified, telemetries is disabled by default.
Default	If the telemetries statement is specified, logging is enabled with the default values for interval and duration.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring PoE on All Interfaces on page 189

template-refresh-rate (Services)

Supported Platforms	SRX Series, vSRX
Syntax	template-refresh-rate;
Hierarchy Level	[edit services flow-monitoring version9 template <i>template-name</i>]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Specify the template refresh rate.
Options	<ul style="list-style-type: none"> • packets—Specify the number of packets. The range is from 1 through 480,000. • seconds—Specify the number of seconds. The range is from 10 through 600.
Required Privilege Level	services—To view this statement in the configuration. services-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Interfaces on page 3

threshold (Interfaces)

Supported Platforms	SRX Series, vSRX
Syntax	threshold <value>;
Hierarchy Level	[edit interfaces <i>interface-name</i> radio-router]
Release Information	Statement introduced in Junos OS Release 10.1.
Description	This option controls the percentage of bandwidth change required for routing updates.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• PPPoE-Based Radio-to-Router Protocols Overview on page 229

traceoptions (Interfaces)

Supported Platforms	SRX Series, vSRX
Syntax	traceoptions
Hierarchy Level	[edit interfaces interface-name traceoptions]
Release Information	Command introduced in Junos OS Release 10.1.
Description	Define tracing operations for individual interfaces. To specify more than one tracing operation, include multiple flag statements.
Options	flag - Tracing parameters
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• PPPoE-Based Radio-to-Router Protocols Overview on page 229

update-server

Supported Platforms	EX Series , J Series , QFX Series , SRX Series
Syntax	update-server;
Hierarchy Level	[edit Interfaces <i>interface-name</i> unit <i>logical-unit-number</i> inet dhcp]
Release Information	Statement introduced in Junos OS Release 8.5 for J Series devices. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 9.2 for SRX Series devices. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Propagate TCP/IP settings learned from an external DHCP server to the DHCP server running on the switch, router, or device.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

vbr rate

Supported Platforms	SRX Series , vSRX
Syntax	vbr rate;
Hierarchy Level	[edit interfaces interface-name atm-options vpi vpi-identifier shaping]
Release Information	Command introduced in Junos OS Release 9.5.
Description	For ATM encapsulation only, define a variable bit rate bandwidth utilization in the traffic-shaping profile.
Options	<ul style="list-style-type: none"> • Burst Size—The maximum burst size that can be sent at the peak rate. • Peak Rate—The maximum instantaneous rate at which the user will transmit. • Sustained Rate—The average rate as measured over a long interval. • CDVT—Cell Delay Variation Tolerance in microseconds (range: 1 – 9999).
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Interfaces on page 3

vdsl-profile

Supported Platforms	SRX320, SRX340, SRX550M
Syntax	vdsl-profile
Hierarchy Level	[edit interfaces interface-name vdsl-options]
Release Information	Command introduced in Junos OS Release 10.1.
Description	Configure the type of VDSL2 profiles. A profile is a table that contains a list of preconfigured VDSL2 settings.
Options	<ul style="list-style-type: none">• Auto (default)• 8a• 8b• 8c• 8d• 12a• 12b• 17a
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• VDSL2 Interface Support on SRX Series Devices on page 77

vendor-id (Interfaces)

Supported Platforms	SRX Series, vSRX
Syntax	vendor-id <i>vendor-id</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family-name</i> dhcp]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Configure a vendor class ID for the Dynamic Host Configuration Protocol (DHCP) client.
Options	<i>vendor-id</i> —vendor class ID.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Interfaces on page 3

vlan-tagging (Interfaces)

Supported Platforms [SRX Series, vSRX](#)

Syntax `vlan-tagging native-vlan-id vlan-id;`

Hierarchy Level [edit interfaces *interface*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure VLAN identifier for untagged packets received on the physical interface of a trunk mode interface.

Options `native-vlan-id`—Configures a VLAN identifier for untagged packets. Enter a number from 0 through 4094.



NOTE: The `native-vlan-id` can be configured only when either `flexible-vlan-tagging mode` or `interface-mode trunk` is configured.

Required Privilege `interface`—To view this statement in the configuration.

Level `interface-control`—To add this statement to the configuration.

Related Documentation

- [Configuring VLAN Tagging on page 49](#)

web-authentication (Interfaces)

Supported Platforms [SRX Series, vSRX](#)

Syntax web-authentication {
 http;
 https;
 redirect-to-https;
}

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

Release Information Statement introduced in Junos OS Release 9.2.
Support for **https** and **redirect-to-https** introduced for high-end SRX Series Services Gateways starting from Junos OS Release 12.1X44-D10 and on vSRX, SRX300, SRX320, SRX340, SRX345, SRX550, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.

Description Enable the Web authentication process for firewall user authentication.

Options **http**—Enable HTTP service.

https—Enable authentication through HTTPS.

redirect-to-https—Redirect Web authentication to HTTPS.

Required Privilege interface—To view this statement in the configuration.

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

Related Documentation

- [Understanding Interfaces on page 3](#)

CHAPTER 19

Operational Commands

- clear oam ethernet connectivity-fault-management path-database
- clear dhcpv6 server binding (Local Server)
- clear ethernet-switching statistics mac-learning
- clear interfaces statistics swfabx
- clear ipv6 neighbors
- clear lacp statistics interfaces
- restart (Reset)
- show chassis fpc (View)
- show chassis hardware (View)
- show ethernet-switching mac-learning-log (View)
- show ethernet-switching table (View)
- show igmp-snooping route (View)
- show interfaces (SRX Series)
- show interfaces diagnostics optics
- show interfaces flow-statistics
- show interfaces queue
- show interfaces statistics (View)
- show interfaces terse zone
- show ipv6 neighbors
- show lacp interfaces (View)
- show lacp statistics interfaces (View)
- show oam ethernet link-fault-management
- show poe controller (View)
- show pppoe interfaces
- show pppoe statistics
- show poe telemetries
- show services accounting
- show services accounting aggregation (View)

- [show services accounting aggregation template \(View\)](#)
- [show services accounting flow-detail \(View\)](#)

clear oam ethernet connectivity-fault-management path-database

Supported Platforms	SRX320, SRX340, SRX345, SRX550M
Syntax	clear oam ethernet connectivity-fault-management path-database maintenance-domain <i>md-name</i> maintenance-association <i>ma-name</i> host < <i>mac-addr</i> >
Release Information	Statement introduced in Junos OS Release 12.1X44-D10.
Description	Clear the relevant path information from the database for the specified remote host.
Options	<p>host—(Optional) MAC address of remote host in xx:xx:xx:xx:xx:xx format.</p> <p>maintenance-association —Name of the maintenance association.</p> <p>maintenance-domain —Name of the maintenance domain.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> <i>show oam ethernet connectivity-fault-management path-database</i>
List of Sample Output	clear oam ethernet connectivity-fault- management path-database on page 319

Sample Output

clear oam ethernet connectivity-fault- management path-database

```
user@host> clear oam ethernet connectivity-fault-management path-database
maintenance-domain private maintenance-association private-ma
Path database entries cleared for the remote-host
```

clear dhcpv6 server binding (Local Server)

Supported Platforms SRX Series

Syntax clear dhcpv6 server binding
<all | *client-id* | *ip-address* | *session-id*>
<interface *interface-name*>
<routing-instance *routing-instance-name*>

Release Information Command introduced in Junos OS Release 10.4.

Description Clear the binding state of a DHCPv6 client from the client table on the DHCPv6 local server.

- Options**
- all—(Optional) Clear the binding state for all DHCPv6 clients.
 - *client-id*—(Optional) Clear the binding state for the DHCPv6 client with the specified client ID (option 1).
 - *ip-address*—(Optional) Clear the binding state for the DHCPv6 client with the specified address.
 - *session-id*—(Optional) Clear the binding state for the DHCPv6 client with the specified session ID.
 - interface *interface-name*—(Optional) Clear the binding state for DHCPv6 clients on the specified interface.
 - routing-instance *routing-instance-name*—(Optional) Clear the binding state for DHCPv6 clients on the specified routing instance.

Required Privilege Level clear

Related Documentation

- *show dhcpv6 server binding (View)*

clear ethernet-switching statistics mac-learning

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, SRX550M, vSRX](#)

Syntax clear ethernet-switching statistics mac-learning

Release Information Command introduced in Junos OS Release 10.1.

Description Clear the media access control (MAC) learning statistics.

- Options**
- **none**—Clear MAC learning statistics on all interfaces.
 - **interface *interface-name***—(Optional) Clear MAC learning statistics on the specified interface.

Required Privilege Level view

Related Documentation

- [show ethernet-switching table \(View\) on page 35](#) show ethernet-switching table

List of Sample Output [clear ethernet-switching statistics mac-learning on page 321](#)
[clear ethernet-switching statistics mac-learning interface interface-name on page 321](#)

Sample Output

clear ethernet-switching statistics mac-learning

```
user@host> clear ethernet-switching statistics mac-learning
```

clear ethernet-switching statistics mac-learning interface interface-name

```
user@host> clear ethernet-switching statistics mac-learning interface interface-name
```

clear interfaces statistics swfabx

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, SRX550M, vSRX](#)

Syntax `clear interfaces statistics <swfab0 | swfab1>`

Release Information Command introduced in Junos OS Release 11.1.

Description Clears interface statistics for the specified swfab interface.

Required Privilege Level clear

Related Documentation

- [show interfaces swfabx](#)

List of Sample Output [clear interfaces statistics <swfab0 | swfab1> on page 322](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

`clear interfaces statistics <swfab0 | swfab1>`

```
user@host> clear interfaces statistics <swfab0 | swfab1>
```

clear ipv6 neighbors

Supported Platforms	SRX1500, SRX320, SRX340, SRX550M, vSRX
Syntax	clear ipv6 neighbors <all host <i>hostname</i> >
Release Information	Command introduced in Junos OS Release 12.1X45-D10.
Description	Clear IPv6 neighbor cache information.
Options	<p>none—Clear all IPv6 neighbor cache information.</p> <p>all—(Optional) Clear all IPv6 neighbor cache information.</p> <p>host <i>hostname</i>—(Optional) Clear the information for the specified IPv6 neighbors.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show ipv6 neighbors on page 404
List of Sample Output	clear ipv6 neighbors on page 323

Sample Output

clear ipv6 neighbors

```
user@host> clear ipv6 neighbors
11:11::2          00:19:e2:4b:61:83  deleted
12:12::2          00:19:e2:4b:61:83  deleted
10:1::2           00:00:0a:00:00:00  deleted
```

clear lacp statistics interfaces

Supported Platforms [SRX Series, vSRX](#)

Syntax `clear lacp statistics interfaces <interface-name>`

Release Information Command modified in Junos OS Release 10.2.

Description Clear the LACP statistics. If you do not specify an interface name, LACP statistics for all interfaces are cleared.

Options *interface-name*—(Optional) Name of an interface.

Required Privilege Level clear

Related Documentation

- [show lacp statistics interfaces \(View\) on page 410](#)
- [Verifying LACP on Redundant Ethernet Interfaces on page 157](#)

Output Fields This command produces no output.

restart (Reset)

Supported Platforms SRX Series, vSRX

Syntax restart
 <application-identification | application-security | audit-process | commitd-service
 | chassis-control | class-of-service | database-replication | datapath-trace-service | ddns
 | dhcp | dhcp-service | dynamic-flow-capture | disk-monitoring | event-processing |
 ethernet-connectivity-fault-management | ethernet-link-fault-management
 | extensible-subscriber-services | fipsd | firewall | firewall-authentication-service
 | general-authentication-service | gracefully | gprs-process | idp-policy | immediately
 | interface-control | ipmi | ipsec-key-management | jflow-service | jnu-management
 | jnx-wmicd-service | jsrp-service | kernel-replication | l2-learning | l2cpd-service | lacp
 | license-service | logical-system-service | mib-process | mountd-service | named-service
 | network-security | network-security-trace | nfsd-service | ntpd-service | pgm
 | pic-services-logging | profilerd | pki-service | remote-operations | rest-api | routing | sampling
 | sampling-route-record | scc-chassisd | secure-neighbor-discovery | security-intelligence
 | security-log | services | service-deployment | simple-mail-client-service | soft | snmp
 | static-routed | statistics-service | subscriber-management | subscriber-management-helper
 | system-log-vital | tunnel-oamd | uac-service | user-ad-authentication | vrrp
 | web-management >

Release Information Command introduced before Junos OS Release 9.2

Description Restart a Junos OS process.



CAUTION: Never restart a software process unless instructed to do so by a customer support engineer. A restart might cause the router to drop calls and interrupt transmission, resulting in possible loss of data.

- Options**
- application-identification—(Optional) Restart the process that identifies an application using intrusion detection and prevention (IDP) to allow or deny traffic based on applications running on standard or nonstandard ports.
 - application-security—(Optional) Restart the application security process.
 - audit-process—(Optional) Restart the RADIUS accounting process that gathers statistical data that can be used for general network monitoring, for analyzing and tracking usage patterns, and for billing a user based upon the amount of time used or the type of services accessed.
 - chassis-control—(Optional) Restart the chassis management process.
 - class-of-service—(Optional) Restart the class-of-service (CoS) process, which controls the router's or switch's CoS configuration.
 - commitd-service—(Optional) Restart the committed services.
 - database-replication—(Optional) Restart the database replication process.
 - datapath-trace-service—(Optional) Restart the Restart the packet path tracing process.

- `ddns`—(Optional) Restart the dynamic domain name system, which dynamically updates IP addresses for registered domain names.
- `dhcp`—(Optional) Restart the software process for a Dynamic Host Configuration Protocol (DHCP) server. A DHCP server allocates network IP addresses and delivers configuration settings to client hosts without user intervention.
- `dhcp-service`—(Optional) Restart the Dynamic Host Configuration Protocol process.
- `disk-monitoring`—(Optional) Restart disk monitoring, which checks the health of the hard disk drive on the Routing Engine.
- `dynamic-flow-capture`—(Optional) Restart the dynamic flow capture (DFC) process, which controls DFC configurations on PIC3 monitoring services cards.
- `ethernet-connectivity-fault-management`—(Optional) Restart the process that provides IEEE 802.1ag Operation, Administration, and Maintenance (OAM) connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.
- `ethernet-link-fault-management`—(Optional) Restart the process that provides the OAM link fault management (LFM) information for Ethernet interfaces.
- `event-processing`—(Optional) Restart the event process (`eventd`).
- `extensible-subscriber-services`—(Optional) Restart the extensible subscriber services process.
- `fipsd`—(Optional) Restart the `fipsd` services.
- `firewall`—(Optional) Restart the firewall management process, which manages the firewall configuration and accepts or rejects packets that are transiting an interface on a router or switch.
- `firewall-authentication-service`—(Optional) Restart the firewall authentication service process.
- `general-authentication-service`—(Optional) Restart the general authentication process.
- `gprs-process`—(Optional) Restart the General Packet Radio Service (GPRS) process.
- `gracefully`—(Optional) Restart the software process.
- `idp-policy`—(Optional) Restart the intrusion detection and prevention (IDP) protocol process.
- `immediately`—(Optional) Immediately restart the software process.
- `interface-control`—(Optional) Restart the interface process, which controls the router's or switch's physical interface devices and logical interfaces.
- `ipmi`—(Optional) Restart the intelligent platform management interface process.
- `ipsec-key-management`—(Optional) Restart the IPsec key management process.
- `jflow-service`—(Optional) Restart `jflow` service process.
- `jnu-management`—(Optional) Restart `jnu` management process.
- `jnx-wmicd-service`—(Optional) Restart `jnx wmicd` service process.

- `jsrp-service`—(Optional) Restart the Juniper Services Redundancy Protocol (jsrdp) process, which controls chassis clustering.
- `kernel-replication`—(Optional) Restart the kernel replication process, which replicates the state of the backup Routing Engine when graceful Routing Engine switchover (GRES) is configured.
- `lACP`—(Optional) Restart the Link Aggregation Control Protocol (LACP) process. LACP provides a standardized means for exchanging information between partner systems on a link. The LACP process allows link aggregation control instances to reach agreement on the identity of the LAG to which a link belongs, moves the link to that LAG, and enables the transmission and reception processes for the link to function in an orderly manner.
- `l2cpd-service`—(High-end SRX Series only) (Optional) Restart the Layer 2 Control Protocol (L2CP) process, which enables features such as L2 protocol tunneling and nonstop bridging.
- `l2-learning`—(Optional) Restart the Layer 2 (L2) address flooding and learning process.
- `license-service`—(Optional) Restart the feature license management process.
- `logical-system-service`—(Optional) Restart the logical system service process.
- `mib-process`—(Optional) Restart the MIB version II process, which provides the router's MIB II agent.
- `mountd-service`—(Optional) Restart the service for Network File System (NFS) mount requests.
- `named-service`—(Optional) Restart the DNS Server process, which is used by a router or a switch to resolve hostnames into addresses.
- `network-security`—(Optional) Restart the network security process.
- `network-security-trace`—(Optional) Restart the network security trace process.
- `nfsd-service`—(Optional) Restart the remote NFS server process, which provides remote file access for applications that need NFS-based transport.
- `ntpd-service`—(Optional) Restart the Network Time Protocol (NTP) process.
- `pgm`—(Optional) Restart the process that implements the Pragmatic General Multicast (PGM) protocol for assisting in the reliable delivery of multicast packets.
- `pic-services-logging`—(Optional) Restart the logging process for some PICs. With this process, also known as `fsad` (the file system access daemon), PICs send special logging information to the Routing Engine for archiving on the hard disk.
- `pki-service`—(Optional) Restart the public key infrastructure (PKI) service process.
- `profillerd`—(Optional) Restart the profiler process.
- `remote-operations`—(Optional) Restart the remote operations process, which provides the ping and traceroute MIBs.
- `rest-api`—(Optional) Restart the rest api process.
- `routing`—(Optional) Restart the routing protocol process (`rpd`).

- `sampling`—(Optional) Restart the sampling process, which performs packet sampling based on particular input interfaces and various fields in the packet header.
- `sampling-route-record`—(Optional) Restart the sampling route record process.
- `scc-chassisd`—(Optional) Restart the scc chassisd process.
- `secure-neighbor-discovery`—(Optional) Restart the secure Neighbor Discovery Protocol (NDP) process, which provides support for protecting NDP messages.
- `security-intelligence`—(Optional) Restart security intelligence process.
- `security-log`—(Optional) Restart the security log process.
- `service-deployment`—(Optional) Restart the service deployment process, which enables Junos OS to work with the Session and Resource Control (SRC) software.
- `services`—(Optional) Restart a service.
- `simple-mail-client-service`—(Optional) Restart the simple mail client service process.
- `snmp`—(Optional) Restart the SNMP process, which enables the monitoring of network devices from a central location and provides the router's or switch's SNMP master agent.
- `static-routed`—(Optional) Restart the static routed process.
- `soft`—(Optional) Reread and reactivate the configuration without completely restarting the software processes. For example, BGP peers stay up and the routing table stays constant. Omitting this option results in a graceful restart of the software process.
- `statistics-service`—(Optional) Restart the process that manages the Packet Forwarding Engine statistics.
- `subscriber-management`—(Optional) Restart the subscriber management process.
- `subscriber-management-helper`—(Optional) Restart the subscriber management helper process.
- `system-log-vital`—(Optional) Restart system log vital process.
- `tunnel-oamd`—(Optional) Restart the tunnel OAM process for L2 tunneled networks.
- `uac-service`—(Optional) Restart the Unified Access Control (UAC) process.
- `user-ad-authentication`—(Optional) Restart User ad Authentication process
- `vrrp`—(Optional) Restart the Virtual Router Redundancy Protocol (VRRP) process, which enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts.
- `web-management`—(Optional) Restart the Web management process.

Required Privilege Level `reset`

Related Documentation • *Restart Commands Overview*

List of Sample Output [restart interfaces on page 329](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

restart interfaces

```
user@host> restart interfaces
interfaces process terminated
interfaces process restarted
```

show chassis fpc (View)

Supported Platforms SRX Series

Syntax `show chassis fpc`
`<detail < fpc-slot > | <node (node-id | local | primary)>> |`
`<node (node-id | local | primary)> |`
`<pic-status < fpc-slot > | <node (node-id | local | primary)>>`

Release Information Command modified in Junos OS Release 9.2.
 Starting with Junos OS Release 15.1X49-D10, the SRX5K-MPC3-100G10G (IOC3) and the SRX5K-MPC3-40G10G (IOC3) are introduced.



NOTE: On SRX5K-MPC3-40G10G (IOC3), all four PICs cannot be powered on. A maximum of two PICs can be powered on at the same time. By default, PIC0 and PIC1 are online.

Use the `set chassis fpc <slot> pic <pic> power off` command to choose the PICs you want to power on.

When you use the `set chassis fpc <slot> pic <pic> power off` command to power off PIC0 and PIC1, PIC2 and PIC3 are automatically turned on.

When you switch from one set of PICs to another set of PICs using the `set chassis fpc <slot> pic <pic> power off` command again, ensure that there is 60 seconds duration between the two actions, otherwise core files are seen during the configuration.

The [Table 21 on page 330](#) summarizes the SRX5K-MPC3-40G10G (IOC3) PICs selected for various configuration scenarios.

Table 21: SRX5K-MPC3-40G10G (IOC3) PIC Selection Summary

CLI Configuration	PIC Selection
Default (i.e. no CLI configuration)	Online: PIC-0, PIC-1 Offline: PIC-2, PIC-3
PIC-1, PIC-2 and PIC-3 powered OFF	Online: PIC-0 Offline: PIC-1, PIC-2, PIC-3
PIC-0, PIC-2 and PIC-3 powered OFF	Online: PIC-1 Offline: PIC-0, PIC-2, PIC-3
PIC-0, PIC-1 and PIC-3 powered OFF	Online: PIC-2 Offline: PIC-0, PIC-1, PIC-3
PIC-0, PIC-1 and PIC-2 powered OFF	Online: PIC-3 Offline: PIC-0, PIC-1, PIC-2

Table 21: SRX5K-MPC3-40G10G (IOC3) PIC Selection Summary (continued)

CLI Configuration	PIC Selection
PIC-2 and PIC-3 powered OFF	Online: PIC-0, PIC-1 Offline: PIC-2, PIC-3
PIC-2 and PIC-3 powered OFF	Online: PIC-0, PIC-1 Offline: PIC-2, PIC-3
PIC-1 and PIC-2 powered OFF	Online: PIC-0, PIC-3 Offline: PIC-1, PIC-2
PIC-0 and PIC-3 powered OFF	Online: PIC-2, PIC-1 Offline: PIC-0, PIC-3
PIC-0 and PIC-1 powered OFF	Online: PIC-2, PIC-3 Offline: PIC-0, PIC-1
All other combinations of PICs being powered OFF (Invalid)	Online: PIC-0, PIC-1 Offline: PIC-2, PIC-3 Default PICs will be selected for the invalid combinations. Also, a system log message will be displayed to indicate the invalid combination PIC selection.

Description Display status information about the installed Flexible PIC Concentrators (FPCs) and PICs.

- Options**
- **none**—Display status information for all FPCs.
 - **detail**—(Optional) Display detailed FPC status information.
 - **fpc-slot** —(Optional) Display information about the FPC in this slot.
 - **node**—(Optional) For chassis cluster configurations, display status information for all FPCs or for the specified FPC on a specific node (device) in the cluster.
 - **node-id** —Identification number of the node. It can be 0 or 1.
 - **local**—Display information about the local node.
 - **primary**—Display information about the primary node.

- **pic-status**—(Optional) Display status information for all FPCs or for the FPC in the specified slot (see *fpc-slot*).

Required Privilege Level view

Related Documentation • [Understanding Interfaces on page 3](#)

List of Sample Output [show chassis fpc on page 333](#)
[show chassis fpc \(SRX5600 and SRX5800 devices\) on page 333](#)
[show chassis fpc \(SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G \(IOC3\) or SRX5K-MPC3-40G10G \(IOC3\) on page 333](#)
[show chassis fpc detail 2 on page 334](#)
[show chassis fpc pic-status \(SRX5600 and SRX5800 devices\) on page 334](#)
[show chassis fpc pic-status \(SRX5600 and SRX5800 devices with SPC2\) on page 334](#)
[show chassis fpc pic-status \(SRX5600 and SRX5800 devices with SRX5K-MPC\) on page 335](#)
[show chassis fpc pic-status \(SRX5600 and SRX5800 devices when Express Path \[formerly known as services offloading\] is configured\) on page 335](#)
[show chassis fpc pic-status \(with 20-Gigabit Ethernet MIC with SFP\) on page 336](#)
[show chassis fpc pic-status\(SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G \(IOC3\) or SRX5K-MPC3-40G10G \(IOC3 and when Express Path \[formerly known as services offloading\] is configured\) on page 336](#)
[show chassis fpc pic-status for HA \(SRX5600 and SRX5800 devices\) on page 336](#)
[show chassis fpc pic-status for HA\(SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G \(IOC3\) or SRX5K-MPC3-40G10G \(IOC3\) on page 337](#)

Output Fields [Table 22 on page 332](#) lists the output fields for the **show chassis fpc** command. Output fields are listed in the approximate order in which they appear.

Table 22: show chassis fpc Output Fields

Field Name	Field Description
Slot or Slot State	Slot number and state. The state can be one of the following conditions: <ul style="list-style-type: none"> • Dead—Held in reset because of errors. • Diag—Slot is being ignored while the device is running diagnostics. • Dormant—Held in reset. • Empty—No FPC is present. • Online—FPC is online and running. • Present—FPC is detected by the device, but is either not supported by the current version of Junos OS or inserted in the wrong slot. The output also states either Hardware Not Supported or Hardware Not In Right Slot. FPC is coming up but not yet online. • Probed—Probe is complete; awaiting restart of the Packet Forwarding Engine (PFE). • Probe-wait—Waiting to be probed.
Temp (C) or Temperature	Temperature of the air passing by the FPC, in degrees Celsius or in both Celsius and Fahrenheit.
Total CPU Utilization (%)	Total percentage of CPU being used by the FPC's processor.

Table 22: show chassis fpc Output Fields (*continued*)

Field Name	Field Description
Interrupt CPU Utilization (%)	Of the total CPU being used by the FPC's processor, the percentage being used for interrupts.
Memory DRAM (MB)	Total DRAM, in megabytes, available to the FPC's processor.
Heap Utilization (%)	Percentage of heap space (dynamic memory) being used by the FPC's processor. If this number exceeds 80 percent, there may be a software problem (memory leak).
Buffer Utilization (%)	Percentage of buffer space being used by the FPC's processor for buffering internal messages.
Start Time	Time when the Routing Engine detected that the FPC was running.
Uptime	How long the Routing Engine has been connected to the FPC and, therefore, how long the FPC has been up and running.
PIC type	(pic-status output only) Type of FPC.

Sample Output

show chassis fpc

```

user@host> show chassis fpc
          Slot State      Temp  CPU Utilization (%)  Memory  Utilization (%)
          0 Online        (C)   Total Interrupt     DRAM (MB) Heap Buffer
          0 Online        ----- CPU less FPC -----
          1 Online        ----- Not Usable -----
          2 Online        ----- CPU less FPC -----

```

show chassis fpc (SRX5600 and SRX5800 devices)

```

user@host> show chassis fpc
          Slot State      Temp  CPU Utilization (%)  Memory  Utilization (%)
          0 Empty        (C)   Total Interrupt     DRAM (MB) Heap Buffer
          1 Empty
          2 Empty
          3 Online        37    3      0      1024    7    42
          4 Empty
          5 Empty
          6 Online        30    8      0      1024    23   30
          7 Empty
          8 Empty
          9 Empty
         10 Empty
         11 Empty

```

show chassis fpc

(SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3))

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%) Total	CPU Utilization (%)		CPU Utilization (%) 1min	CPU Utilization (%) 5min	CPU Utilization (%) 15min	Memory DRAM (MB)
				Interrupt	Heap				
				Buffer					
0	Online	36	20	0	20	19	19	1024	
1	Online	35	8	0	8	8	8	2048	
2	Online	40	21	0	14	20	20	3584	
				5	13				

Sample Output

show chassis fpc detail 2

```

user@host> show chassis fpc detail 2
Slot 2 information:
State                               Online
Temperature                          37
Total CPU DRAM                       1024 MB
Total RLDRAM                          0 MB
Total DDR DRAM                        0 MB
Start time:                           2012-07-18 07:18:50 PDT
Uptime:                               4 days, 21 hours, 51 minutes, 59 seconds

Max Power Consumption                 0 Watts
    
```

Sample Output

show chassis fpc pic-status (SRX5600 and SRX5800 devices)

```

user@host> show chassis fpc pic-status
Slot 3  Online  SRX5k SPC
PIC 0  Online  SPU Cp
PIC 1  Online  SPU Flow
Slot 6  Online  SRX5k DPC 4x 10GE
PIC 0  Online  1x 10GE(LAN/WAN) RichQ
PIC 1  Online  1x 10GE(LAN/WAN) RichQ
PIC 2  Online  1x 10GE(LAN/WAN) RichQ
PIC 3  Online  1x 10GE(LAN/WAN) RichQ
    
```

show chassis fpc pic-status (SRX5600 and SRX5800 devices with SPC2)

```

user@host> show chassis fpc pic-status

Slot 0  Online  SRX5k DPC 40x 1GE
PIC 0  Online  10x 1GE RichQ
PIC 1  Online  10x 1GE RichQ
PIC 2  Online  10x 1GE RichQ
PIC 3  Online  10x 1GE RichQ
Slot 2  Online  SRX5k SPC II
PIC 0  Online  SPU Cp
PIC 1  Online  SPU Flow
PIC 2  Online  SPU Flow
PIC 3  Online  SPU Flow
Slot 3  Online  SRX5k SPC II
    
```

```

PIC 0 Online      SPU Flow
PIC 1 Online      SPU Flow
PIC 2 Online      SPU Flow
PIC 3 Online      SPU Flow
Slot 5 Online     SRX5k SPC
PIC 0 Online      SPU Flow
PIC 1 Online      SPU Flow

```

show chassis fpc pic-status (SRX5600 and SRX5800 devices with SRX5K-MPC)

```

user@host> show chassis fpc pic-status

Slot 0 Online      SRX5k SPC II
  PIC 0 Online     SPU Cp
  PIC 1 Online     SPU Flow
  PIC 2 Online     SPU Flow
  PIC 3 Online     SPU Flow
Slot 1 Online      SRX5k SPC II
  PIC 0 Online     SPU Flow
  PIC 1 Online     SPU Flow
  PIC 2 Online     SPU Flow
  PIC 3 Online     SPU Flow
Slot 2 Online      SRX5k DPC 4X 10GE
  PIC 0 Online     1x 10GE(LAN/WAN) RichQ
  PIC 1 Online     1x 10GE(LAN/WAN) RichQ
  PIC 2 Online     1x 10GE(LAN/WAN) RichQ
  PIC 3 Online     1x 10GE(LAN/WAN) RichQ
Slot 6 Offline     SRX5k SPC II
Slot 9 Online      SRX5k SPC II
  PIC 0 Online     SPU Flow
  PIC 1 Online     SPU Flow
  PIC 2 Online     SPU Flow
  PIC 3 Online     SPU Flow
Slot 10 Online     SRX5k IOC II
  PIC 0 Online     10x 10GE SFP+
  PIC 2 Online     1x 100GE CFP
Slot 11 Online     SRX5k IOC II
  PIC 0 Online     1x 100GE CFP
  PIC 2 Online     2x 40GE QSFP+

```

show chassis fpc pic-status (SRX5600 and SRX5800 devices when Express Path [formerly known as services offloading] is configured)

```

user@host> show chassis fpc pic-status

Slot 0 Offline     SRX5k DPC 40x 1GE
Slot 1 Online      SRX5k SPC II
  PIC 0 Online     SPU Cp
  PIC 1 Online     SPU Flow
  PIC 2 Online     SPU Flow
  PIC 3 Online     SPU Flow
Slot 2 Offline     SRX5k SPC
Slot 4 Online      SRX5k IOC3 24XGE+6XLG
  PIC 2 Online     3x 40GE QSFP+- np-cache/services-offload
  PIC 3 Online     3x 40GE QSFP+- np-cache/services-offload
Slot 5 Online      SRX5k IOC II
  PIC 0 Online     10x 1GE(LAN) SFP- np-cache/services-offload
  PIC 1 Online     10x 1GE(LAN) SFP- np-cache/services-offload
  PIC 2 Online     10x 10GE SFP+- np-cache/services-offload

```

show chassis fpc pic-status (with 20-Gigabit Ethernet MIC with SFP)

```
user@host> show chassis fpc pic-status
```

```
node0:
```

```
-----
Slot 0  Online      SRX5k SPC II
  PIC 0  Online      SPU Cp
  PIC 1  Online      SPU Flow
  PIC 2  Online      SPU Flow
  PIC 3  Online      SPU Flow
Slot 1  Offline     SRX5k SPC II
Slot 2  Online      SRX5k DPC 4X 10GE
  PIC 0  Online      1x 10GE(LAN/WAN) RichQ
  PIC 1  Online      1x 10GE(LAN/WAN) RichQ
  PIC 2  Online      1x 10GE(LAN/WAN) RichQ
  PIC 3  Online      1x 10GE(LAN/WAN) RichQ
Slot 9  Online      SRX5k IOC II
  PIC 0  Online      10x 1GE(LAN) SFP
  PIC 1  Online      10x 1GE(LAN) SFP
  PIC 2  Online      10x 1GE(LAN) SFP
  PIC 3  Online      10x 1GE(LAN) SFP
Slot 10 Online      SRX5k IOC II
  PIC 0  Online      10x 10GE SFP+
  PIC 2  Online      1x 100GE CFP
Slot 11 Offline     SRX5k IOC II
```

show chassis fpc pic-status

(SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3 and when Express Path [formerly known as services offloading] is configured)

```
user@host> show chassis fpc pic-status
```

```
Slot 0  Offline     SRX5k DPC 40x 1GE
Slot 1  Online      SRX5k SPC II
  PIC 0  Online      SPU Cp
  PIC 1  Online      SPU Flow
  PIC 2  Online      SPU Flow
  PIC 3  Online      SPU Flow
Slot 2  Offline     SRX5k SPC
Slot 4  Online      SRX5k IOC3 24XGE+6XLG
  PIC 2  Online      3x 40GE QSFP+- np-cache/services-offload
  PIC 3  Online      3x 40GE QSFP+- np-cache/services-offload
Slot 5  Online      SRX5k IOC II
  PIC 0  Online      10x 1GE(LAN) SFP- np-cache/services-offload
  PIC 1  Online      10x 1GE(LAN) SFP- np-cache/services-offload
  PIC 2  Online      10x 10GE SFP+- np-cache/services-offload
```

Sample Output

show chassis fpc pic-status for HA (SRX5600 and SRX5800 devices)

```
user@host> show chassis fpc pic-status
```

```
node0:
```

```
-----
Slot 4  Online      SRX5k DPC 40x 1GE
  PIC 0  Online      10x 1GE RichQ
  PIC 1  Online      10x 1GE RichQ
  PIC 2  Online      10x 1GE RichQ
  PIC 3  Online      10x 1GE RichQ
Slot 5  Online      SRX5k SPC
```

```

PIC 0 Online      SPU Cp-Flow
PIC 1 Online      SPU Flow

```

```
node1:
```

```

-----
Slot 4 Online      SRX5k DPC 40x 1GE
PIC 0 Online      10x 1GE RichQ
PIC 1 Online      10x 1GE RichQ
PIC 2 Online      10x 1GE RichQ
PIC 3 Online      10x 1GE RichQ
Slot 5 Online      SRX5k SPC
PIC 0 Online      SPU Cp-Flow
PIC 1 Online      SPU Flow

```

show chassis fpc pic-status for HA
(SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3))

```

user@host> show chassis fpc pic-status
user@host> show chassis fpc pic-status
node0:

```

```

-----
Slot 2 Online      SRX5k IOC3 24XGE+6XLG
PIC 0 Online      12x 10GE SFP+
PIC 1 Online      12x 10GE SFP+
PIC 2 Offline     3x 40GE QSFP+
PIC 3 Offline     3x 40GE QSFP+
Slot 4 Online      SRX5k IOC II
PIC 2 Online      10x 10GE SFP+
Slot 5 Online      SRX5k SPC II
PIC 0 Online      SPU Cp
PIC 1 Online      SPU Flow
PIC 2 Offline
PIC 3 Offline

```

```
node1:
```

```

-----
Slot 2 Online      SRX5k IOC3 24XGE+6XLG
PIC 0 Online      12x 10GE SFP+
PIC 1 Online      12x 10GE SFP+
PIC 2 Offline     3x 40GE QSFP+
PIC 3 Offline     3x 40GE QSFP+
Slot 4 Online      SRX5k IOC II
PIC 2 Online      10x 10GE SFP+
Slot 5 Online      SRX5k SPC II
PIC 0 Online      SPU Cp
PIC 1 Online      SPU Flow
PIC 2 Offline
PIC 3 Offline

```

show chassis hardware (View)

Supported Platforms	SRX Series
Syntax	<code>show chassis hardware</code> <code><clei-models detail extensive models node (<i>node-id</i> all local primary)></code>
Release Information	Command introduced in Junos OS Release 9.2. Command modified in Junos OS Release 9.2 to include node option.
Description	Display chassis hardware information.
Options	<ul style="list-style-type: none"> • clei-models—(Optional) Display Common Language Equipment Identifier Code (CLEI) barcode and model number for orderable field-replaceable units (FRUs). • detail extensive—(Optional) Display the specified level of output. • models—(Optional) Display model numbers and part numbers for orderable FRUs. • node—(Optional) For chassis cluster configurations, display chassis hardware information on a specific node (device) in the cluster. <ul style="list-style-type: none"> • node-id—Identification number of the node. It can be 0 or 1. • local—Display information about the local node. • primary—Display information about the primary node.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Juniper Networks Devices Processing Overview • Interface Naming Conventions on page 8
Output Fields	Table 23 on page 338 lists the output fields for the show chassis hardware command. Output fields are listed in the approximate order in which they appear.

Table 23: show chassis hardware Output Fields

Field Name	Field Description
Item	Chassis component—Information about the backplane; power supplies; fan trays; Routing Engine; each Physical Interface Module (PIM)—reported as FPC and PIC—and each fan, blower, and impeller.
Version	Revision level of the chassis component.
Part Number	Part number for the chassis component.
Serial Number	Serial number of the chassis component. The serial number of the backplane is also the serial number of the device chassis. Use this serial number when you need to contact Juniper Networks Customer Support about the device chassis.

Table 23: show chassis hardware Output Fields (*continued*)

Field Name	Field Description
Assb ID or Assembly ID	Identification number that describes the FRU hardware.
FRU model number	Model number of FRU hardware component.
CLEI code	Common Language Equipment Identifier code. This value is displayed only for hardware components that use ID EEPROM format v2. This value is not displayed for components that use ID EEPROM format v1.
EEPROM Version	ID EEPROM version used by hardware component: 0x01 (version 1) or 0x02 (version 2).

Table 23: show chassis hardware Output Fields (*continued*)

Field Name	Field Description
Description	<p>Brief description of the hardware item:</p> <ul style="list-style-type: none"> • Type of power supply. • Switch Control Board (SCB) <p>Starting with Junos OS Release 12.1X47-D15, the SRX5K-SCBE (SCB2) is introduced.</p> <ul style="list-style-type: none"> • There are three SCB slots in SRX5800 devices. The third slot can be used for an SCB or an FPC. When an SRX5K-SCB was used , the third SCB slot was used as an FPC. SCB redundancy is provided in chassis cluster mode. • With an SCB2, a third SCB is supported. If a third SCB is plugged in, it provides intra-chassis fabric redundancy. • The Ethernet switch in the SCB2 provides the Ethernet connectivity among all the FPCs and the Routing Engine. The Routing Engine uses this connectivity to distribute forwarding and routing tables to the FPCs. The FPCs use this connectivity to send exception packets to the Routing Engine. • Fabric connects all FPCs in the data plane. The Fabric Manager executes on the Routing Engine and controls the fabric system in the chassis. Packet Forwarding Engines on the FPC and fabric planes on the SCB are connected through HSL2 channels. • SCB2 supports HSL2 with both 3.11 Gbps and 6.22 Gbps (SerDes) link speed and various HSL2 modes. When an FPC is brought online, the link speed and HSL2 mode are determined by the type of FPC. <p>Starting with Junos OS Release 15.1X49-D10, the SRX5K-SCB3 (SCB3) with enhanced midplanes is introduced.</p> <ul style="list-style-type: none"> • All existing SCB software that is supported by SCB2 is supported on SCB3. • SRX5K-RE-1800X4 (RE2). Mixed Routing Engine use is not supported. • SCB3 works with the SRX5K-MPC (IOC2), SRX5K-MPC3-100G10G (IOC3), SRX5K-MPC3-40G10G (IOC3), and SRX5K-SPC-4-15-320 (SPC2) with current midplanes and the new enhanced midplanes. • Mixed SCB use is not supported. If an SCB2 and an SCB3 are used, the system will only power on the master Routing Engine's SCB and will power off the other SCBs. Only the SCB in slot 0 is powered on and a system log is generated. • SCB3 supports up to 400 Gbps per slot with old midplanes and up to 500 Gbps per slot with new midplanes. • SCB3 supports fabric intra-chassis redundancy. • SCB3 supports the same chassis cluster function as the SRX5K-SCB (SCB1) and the SRX5K-SCBE (SCB2), except for in-service software upgrade (ISSU) and in-service hardware upgrade (ISHU). • SCB3 has a second external Ethernet port. • Fabric bandwidth increasing mode is not supported.

Table 23: show chassis hardware Output Fields (*continued*)

Field Name	Field Description
	<ul style="list-style-type: none"> • Type of Flexible PIC Concentrator (FPC), Physical Interface Card (PIC), Modular Interface Cards (MICs), and PIMs. • IOCs <p>Starting with Junos OS Release 15.1X49-D10, the SRX5K-MPC3-100G10G (IOC3) and the SRX5K-MPC3-40G10G (IOC3) are introduced.</p> <ul style="list-style-type: none"> • IOC3 has two types of IOC3 MPCs, which have different built-in MICs: the 24x10GE + 6x40GE MPC and the 2x100GE + 4x10GE MPC. • IOC3 supports SCB3 and SRX5000 line backplane and enhanced backplane. • IOC3 can only work with SRX5000 line SCB2 and SCB3. If an SRX5000 line SCB is detected, IOC3 is offline, an FPC misconfiguration alarm is raised, and a system log message is generated. • IOC3 interoperates with SCB2 and SCB3. • IOC3 interoperates with the SRX5K-SPC-4-15-320 (SPC2) and the SRX5K-MPC (IOC2). • The maximum power consumption for one IOC3 is 645W. An enhanced power module must be used. • The IOC3 does not support the following command to set a PIC to go offline or online: request chassis pic fpc-slot <fpc-slot> pic-slot <pic-slot> <offline online> . • IOC3 supports 240 Gbps of throughput with the enhanced SRX5000 line backplane. • Chassis cluster functions the same as for the SRX5000 line IOC2. • IOC3 supports intra-chassis and inter-chassis fabric redundancy mode. • IOC3 supports ISSU and ISHU in chassis cluster mode. • IOC3 supports intra-FPC and Inter-FPC Express Path (previously known as <i>services offloading</i>) with IPv4. • NAT of IPv4 and IPv6 in normal mode and IPv4 for Express Path mode. • All four PICs on the 24x10GE + 6x40GE cannot be powered on. A maximum of two PICs can be powered on at the same time. Use the set chassis fpc <slot> pic <pic> power off command to choose the PICs you want to power on. <p>NOTE: Fabric bandwidth increasing mode is not supported on IOC3.</p> • SRX Clustering Module (SCM) • Fan tray • For hosts, the Routing Engine type. <ul style="list-style-type: none"> • Starting with Junos OS Release 12.1X47-D15, the SRX5K-RE-1800X4 (RE2) Routing Engine is introduced. • The RE2 has an Intel Quad core Xeon processor, 16 GB of DRAM, and a 128-GB solid-state drive (SSD). The number 1800 refers to the speed of the processor (1.8 GHz). The maximum required power for this Routing Engine is 90W. <p>NOTE: The RE2 provides significantly better performance than the previously used Routing Engine, even with a single core.</p>

show chassis hardware

show chassis hardware

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               CM0715AK0021  SRX1500
Midplane      REV 08   750-058562  ACMA4255       SRX1500
CB 0          REV 08   711-053838  ACMA7529       CPU Board SRX700E
Routing Engine 0
FPC 0        REV 07   711-053832  ACMA3311       SRX Routing Engine
PIC 0        BUILTIN BUILTIN       FEB
Xcvr 12      REV 01   740-014132  61521013      12x1G-T-4x1G-SFP-4x10G
Xcvr 13      REV 02   740-013111  A281604       SFP-T
Xcvr 14      REV 02   740-011613  NRN30NV       SFP-SX
Xcvr 15      REV 02   740-011613  NRN2PWV       SFP-SX
Xcvr 16      REV 01   740-021308  AJA17B5       SFP+-10G-SR
Xcvr 17      REV 01   740-021308  MSP056B       SFP+-10G-SR
Xcvr 18      REV 01   740-031980  AS920WJ       SFP+-10G-SR
Xcvr 19      REV 01   740-031980  AS92W5N       SFP+-10G-SR
Power Supply 0 REV 01   740-055217  1EDP42500JZ   PS 400W 90-264V AC in
Fan Tray 0
Airflow - AFO
Fan Tray 1
Airflow - AFO
Fan Tray 2
Airflow - AFO
Fan Tray 3
Airflow - AFO
SRX1500 0, Front to Back
SRX1500 1, Front to Back
SRX1500 2, Front to Back
SRX1500 3, Front to Back

```

show chassis hardware (SRX5600 and SRX5800 devices for SRX5K-MPC)

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN12170EAAGA  SRX 5800
Midplane      REV 01   710-041799  ACAX3849       SRX 5800 Backplane
FPM Board     REV 01   710-024632  CAAX7297       Front Panel Display
PDM           Rev 03   740-013110  QCS170250DU   Power Distribution Module
PEM 0        Rev 03   740-034724  QCS17020203F  PS 4.1kW; 200-240V AC in
PEM 1        Rev 03   740-034724  QCS17020203C  PS 4.1kW; 200-240V AC in
PEM 2        Rev 04   740-034724  QCS17100200A  PS 4.1kW; 200-240V AC in
PEM 3        Rev 03   740-034724  QCS17080200M  PS 4.1kW; 200-240V AC in
Routing Engine 0 REV 11   740-023530  9012047437    SRX5k RE-13-20
CB 0         REV 09   710-024802  CAAX7202      SRX5k SCB
CB 1         REV 09   710-024802  CAAX7157      SRX5k SCB
FPC 0        REV 07   750-044175  CAAD0791      SRX5k SPC II
CPU          BUILTIN BUILTIN       SRX5k DPC PPC
PIC 0        BUILTIN BUILTIN       SPU Cp
PIC 1        BUILTIN BUILTIN       SPU Flow
PIC 2        BUILTIN BUILTIN       SPU Flow
PIC 3        BUILTIN BUILTIN       SPU Flow
FPC 1        REV 07   750-044175  CAAD0751      SRX5k SPC II
CPU          BUILTIN BUILTIN       SRX5k DPC PPC
PIC 0        BUILTIN BUILTIN       SPU Flow

```

PIC 1		BUILTIN	BUILTIN	SPU Flow
PIC 2		BUILTIN	BUILTIN	SPU Flow
PIC 3		BUILTIN	BUILTIN	SPU Flow
FPC 2	REV 28	750-020751	CAAW1817	SRX5k DPC 4X 10GE
CPU	REV 04	710-024633	CAAZ5269	SRX5k DPC PMB
PIC 0		BUILTIN	BUILTIN	1x 10GE(LAN/WAN) RichQ
Xcvr 0	REV 02	740-014289	T10A00404	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	1x 10GE(LAN/WAN) RichQ
PIC 2		BUILTIN	BUILTIN	1x 10GE(LAN/WAN) RichQ
PIC 3		BUILTIN	BUILTIN	1x 10GE(LAN/WAN) RichQ
FPC 6	REV 02	750-044175	ZY2552	SRX5k SPC II
CPU		BUILTIN	BUILTIN	SRX5k DPC PPC
FPC 9	REV 10	750-044175	CAAP5932	SRX5k SPC II
CPU		BUILTIN	BUILTIN	SRX5k DPC PPC
PIC 0		BUILTIN	BUILTIN	SPU Flow
PIC 1		BUILTIN	BUILTIN	SPU Flow
PIC 2		BUILTIN	BUILTIN	SPU Flow
PIC 3		BUILTIN	BUILTIN	SPU Flow
FPC 10	REV 22	750-043157	ZH8192	SRX5k IOC II CPU
REV 08	711-043360	YX3879		SRX5k MPC PMB
MIC 0	REV 01	750-049488	YZ2084	10x 10GE SFP+
PIC 0		BUILTIN	BUILTIN	10x 10GE SFP+
Xcvr 0	REV 01	740-031980	AMBOHG3	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AM20B6F	SFP+-10G-SR
MIC 1	REV 19	750-049486	CAAH3504	1x 100GE CFP
PIC 2		BUILTIN	BUILTIN	1x 100GE CFP
Xcvr 0	REV 01	740-035329	X000D375	CFP-100G-SR10
FPC 11	REV 07.04.07	750-043157	CAAJ8771	SRX5k IOC II CPU
REV 08	711-043360	CAAJ3881		SRX5k MPC PMB
MIC 0	REV 19	750-049486	CAAH0979	1x 100GE CFP
PIC 0		BUILTIN	BUILTIN	1x 100GE CFP
Xcvr 0	REV 01	740-035329	UP1020Z	CFP-100G-SR10
MIC 1	REV 08	750-049487	CAAM1160	2x 40GE QSFP+
PIC 2		BUILTIN	BUILTIN	2x 40GE QSFP+
Xcvr 0	REV 01	740-032986	QB151094	QSFP+-40G-SR4
Xcvr 1	REV 01	740-032986	QB160509	QSFP+-40G-SR4
Fan Tray 0	REV 04	740-035409	ACAE0875	Enhanced Fan Tray
Fan Tray 1	REV 04	740-035409	ACAE0876	Enhanced Fan Tray

show chassis hardware (with 20-Gigabit Ethernet MIC with SFP)

```
user@host> show chassis hardware
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN108DA5AAGA	SRX 5800
Midplane	REV 02	710-013698	TR0037	SRX 5600 Midplane
FPM Board	REV 02	710-014974	JY4635	Front Panel Display
PDM	Rev 02	740-013110	QCS10465005	Power Distribution Module
PEM 0	Rev 03	740-023514	QCS11154040	PS 1.7kW; 200-240VAC in
PEM 2	Rev 02	740-023514	QCS10504014	PS 1.7kW; 200-240VAC in
Routing Engine 0	REV 05	740-015113	1000681023	RE-S-1300
CB 0	REV 05	710-013385	JY4775	SRX5k SCB
FPC 1	REV 17	750-020751	WZ6349	SRX5k DPC 4X 10GE
CPU	REV 02	710-024633	WZ0718	SRX5k DPC PMB
PIC 0		BUILTIN	BUILTIN	1x 10GE(LAN/WAN) RichQ
Xcvr 0		NON-JNPR	C724XM088	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	1x 10GE(LAN/WAN) RichQ
Xcvr 0	REV 02	740-011571	C831XJ08S	XFP-10G-SR
PIC 2		BUILTIN	BUILTIN	1x 10GE(LAN/WAN) RichQ
PIC 3		BUILTIN	BUILTIN	1x 10GE(LAN/WAN) RichQ
FPC 3	REV 22	750-043157	ZH8189	SRX5k IOC II

CPU	REV 06	711-043360	YX3912	SRX5k MPC PMB
MIC 0	REV 01	750-055732	CACF9115	20x 1GE(LAN) SFP
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 2	REV 02	740-013111	B358549	SFP-T
Xcvr 9	REV 02	740-011613	PNB1FQS	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 9	REV 02	740-011613	PNB1FFF	SFP-SX
FPC 5	REV 01	750-027945	JW9665	SRX5k FIOC
CPU				
FPC 8	REV 08	750-023996	XA7234	SRX5k SPC
CPU	REV 02	710-024633	XA1599	SRX5k DPC PMB
PIC 0		BUILTIN	BUILTIN	SPU Cp-Flow
PIC 1		BUILTIN	BUILTIN	SPU Flow
Fan Tray 0	REV 03	740-014971	TP0902	Fan Tray
Fan Tray 1	REV 01	740-014971	TP0121	Fan Tray

show chassis hardware

(SRX5600 and SRX5800 devices with SRX5000 line SRX5K-SCBE [SCB2] and SRX5K-RE-1800X4 [RE2])

user@host> show chassis hardware

node0:

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN1251EA1AGB	SRX5600
Midplane	REV 01	760-063936	ACRE2657	Enhanced SRX5600 Midplane
FPM Board	REV 01	710-024631	CABY3551	Front Panel Display
PEM 0	Rev 03	740-034701	QCS13380901P	PS 1.4-2.6kW; 90-264V
AC in				
PEM 1	Rev 03	740-034701	QCS133809019	PS 1.4-2.6kW; 90-264V
AC in				
Routing Engine 0	REV 02	740-056658	9009210105	SRX5k RE-1800X4
Routing Engine 1	REV 02	740-056658	9013115551	SRX5k RE-1800X4
CB 0	REV 01	750-062257	CADW3663	SRX5k SCB3
CB 1	REV 01	750-062257	CADZ3263	SRX5k SCB3
FPC 0	REV 18	750-054877	CABG6043	SRX5k SPC II
CPU		BUILTIN	BUILTIN	SRX5k DPC PPC
PIC 0		BUILTIN	BUILTIN	SPU Cp
PIC 1		BUILTIN	BUILTIN	SPU Flow
PIC 2		BUILTIN	BUILTIN	SPU Flow
PIC 3		BUILTIN	BUILTIN	SPU Flow
FPC 1	REV 01	750-062243	CAEE5918	SRX5k IOC3 24XGE+6XLG
CPU	REV 02	711-062244	CADX8509	RMPC PMB
PIC 0		BUILTIN	BUILTIN	12x 10GE SFP+
Xcvr 0	REV 01	740-031980	273363A01891	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	273363A01915	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	ANA0BK6	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AP407GA	SFP+-10G-SR
Xcvr 9	REV 01	740-021308	MUC20G1	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	12x 10GE SFP+
PIC 2		BUILTIN	BUILTIN	3x 40GE QSFP+
PIC 3		BUILTIN	BUILTIN	3x 40GE QSFP+
WAN MEZZ	REV 15	750-049136	CAEE5845	MPC5E 24XGE OTN Mezz
FPC 3	REV 11	750-043157	CACL7452	SRX5k IOC II
CPU	REV 04	711-043360	CACP1977	SRX5k MPC PMB
MIC 0	REV 04	750-049488	CABL4759	10x 10GE SFP+
PIC 0		BUILTIN	BUILTIN	10x 10GE SFP+
Xcvr 0	REV 01	740-021308	CF36KMOSY	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	MUCOMF2	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	CF36KM01S	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	MUC229N	SFP+-10G-SR

```

FPC 5          REV 07  750-044175  CAAD0764          SRX5k SPC II
CPU            BUILTIN  BUILTIN          SRX5k DPC PPC
PIC 0         BUILTIN  BUILTIN          SPU Flow
PIC 1         BUILTIN  BUILTIN          SPU Flow
PIC 2         BUILTIN  BUILTIN          SPU Flow
PIC 3         BUILTIN  BUILTIN          SPU Flow
Fan Tray      Enhanced Fan Tray

```

```
node1:
```

```

-----
Hardware inventory:
Item             Version  Part number  Serial number  Description
Chassis                               JN124FE77AGB  SRX5600
Midplane        REV 01  760-063936  ACRE2970      Enhanced SRX5600 Midplane
FPM Board       REV 01  710-024631  CABY3552      Front Panel Display
PEM 0           Rev 03  740-034701  QCS133809028 PS 1.4-2.6kW; 90-264V
AC in
PEM 1           Rev 03  740-034701  QCS133809027 PS 1.4-2.6kW; 90-264V
AC in
Routing Engine 0 REV 02  740-056658  9009218294    SRX5k RE-1800X4
Routing Engine 1 REV 02  740-056658  9013104758    SRX5k RE-1800X4
CB 0            REV 01  750-062257  CAEB8180      SRX5k SCB3
CB 1            REV 01  750-062257  CADZ3334      SRX5k SCB3
FPC 0           REV 18  750-054877  CACJ9834      SRX5k SPC II
CPU            BUILTIN  BUILTIN          SRX5k DPC PPC
PIC 0         BUILTIN  BUILTIN          SPU Cp
PIC 1         BUILTIN  BUILTIN          SPU Flow
PIC 2         BUILTIN  BUILTIN          SPU Flow
PIC 3         BUILTIN  BUILTIN          SPU Flow
FPC 1           REV 01  750-062243  CAEB0981      SRX5k IOC3 24XGE+6XLG
CPU            REV 02  711-062244  CAEA4644      RMPC PMB
PIC 0         BUILTIN  BUILTIN          12x 10GE SFP+
  Xcvr 0       REV 01  740-031980  AP41BLH       SFP+-10G-SR
  Xcvr 1       REV 01  740-031980  AQ400SL       SFP+-10G-SR
  Xcvr 2       REV 01  740-031980  AP422LJ       SFP+-10G-SR
  Xcvr 3       REV 01  740-021308  AMGORBT       SFP+-10G-SR
  Xcvr 9       REV 01  740-021308  MUC2FRG       SFP+-10G-SR
PIC 1         BUILTIN  BUILTIN          12x 10GE SFP+
PIC 2         BUILTIN  BUILTIN          3x 40GE QSFP+
PIC 3         BUILTIN  BUILTIN          3x 40GE QSFP+
WAN MEZZ       REV 15  750-049136  CAEA4837      MPC5E 24XGE OTN Mezz
FPC 3           REV 11  750-043157  CACA8784      SRX5k IOC II
CPU            REV 04  711-043360  CACA8820      SRX5k MPC PMB
MIC 0           REV 05  750-049488  CADF0521      10x 10GE SFP+
PIC 0         BUILTIN  BUILTIN          10x 10GE SFP+
  Xcvr 0       REV 01  740-030658  AD1130A00PV   SFP+-10G-USR
  Xcvr 1       REV 01  740-031980  AN40MVV       SFP+-10G-SR
  Xcvr 2       REV 01  740-021308  CF36KM37B     SFP+-10G-SR
  Xcvr 3       REV 01  740-021308  AD153830DSZ   SFP+-10G-SR
MIC 1           REV 01  750-049487  CABB5961      2x 40GE QSFP+
PIC 2         BUILTIN  BUILTIN          2x 40GE QSFP+
  Xcvr 1       REV 01  740-032986  QB160513      QSFP+-40G-SR4
FPC 5           REV 02  750-044175  ZY2569        SRX5k SPC II
CPU            BUILTIN  BUILTIN          SRX5k DPC PPC
PIC 0         BUILTIN  BUILTIN          SPU Flow
PIC 1         BUILTIN  BUILTIN          SPU Flow
PIC 2         BUILTIN  BUILTIN          SPU Flow
PIC 3         BUILTIN  BUILTIN          SPU Flow
Fan Tray      Enhanced Fan Tray

```

```
show chassis hardware
```

(SRX5400, SRX5600, and SRX5800 devices with SRX5000 line SRX5K-SCB3 [SCB3] with enhanced midplanes and SRX5K-MPC3-100G10G [IOC3] or SRX5K-MPC3-40G10G [IOC3])

```

user@host> show chassis hardware
node0:
-----
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN1250870AGB  SRX5600
Midplane      REV 01   760-063936  ACRE2578      Enhanced SRX5600 Midplane

FPM Board
PEM 0         Rev 03   740-034701  QCS13090900T  PS 1.4-2.6kW; 90-264V A

PEM 1         Rev 03   740-034701  QCS13090904T  PS 1.4-2.6kW; 90-264V A

Routing Engine 0 REV 01   740-056658  9009196496    SRX5k RE-1800X4
CB 0          REV 01   750-062257  CAEC2501      SRX5k SCB3
FPC 0         REV 10   750-056758  CADC8067      SRX5k SPC II
  CPU         BUILTIN BUILTIN      BUILTIN      SRX5k DPC PPC
  PIC 0       BUILTIN BUILTIN      BUILTIN      SPU Cp
  PIC 1       BUILTIN BUILTIN      BUILTIN      SPU Flow
  PIC 2       BUILTIN BUILTIN      BUILTIN      SPU Flow
  PIC 3       BUILTIN BUILTIN      BUILTIN      SPU Flow
FPC 2         REV 01   750-062243  CAEE5924      SRX5k IOC3 24XGE+6XLG
  CPU         REV 01   711-062244  CAEB4890      SRX5k IOC3 PMB
  PIC 0       BUILTIN BUILTIN      BUILTIN      12x 10GE SFP+
  PIC 1       BUILTIN BUILTIN      BUILTIN      12x 10GE SFP+
  PIC 2       BUILTIN BUILTIN      BUILTIN      3x 40GE QSFP+
  Xcvr 0      REV 01   740-038623  MOC13156230449 QSFP+-40G-CU1M
  Xcvr 2      REV 01   740-038623  MOC13156230449 QSFP+-40G-CU1M
  PIC 3       BUILTIN BUILTIN      BUILTIN      3x 40GE QSFP+
WAN MEZZ      REV 01   750-062682  CAEE5817      24x 10GE SFP+ Mezz
FPC 4         REV 11   750-043157  CACY1595      SRX5k IOC II
  CPU         REV 04   711-043360  CACZ8879      SRX5k MPC PMB
  MIC 1       REV 04   750-049488  CACM6062      10x 10GE SFP+
  PIC 2       BUILTIN BUILTIN      BUILTIN      10x 10GE SFP+
  Xcvr 7      REV 01   740-021308  AD1439301TU   SFP+-10G-SR
  Xcvr 8      REV 01   740-021308  AD1439301SD   SFP+-10G-SR
  Xcvr 9      REV 01   740-021308  AD1439301TS   SFP+-10G-SR
FPC 5         REV 05   750-044175  ZZ1371        SRX5k SPC II
  CPU         BUILTIN BUILTIN      BUILTIN      SRX5k DPC PPC
  PIC 0       BUILTIN BUILTIN      BUILTIN      SPU Flow
  PIC 1       BUILTIN BUILTIN      BUILTIN      SPU Flow
  PIC 2       BUILTIN BUILTIN      BUILTIN      SPU Flow
  PIC 3       BUILTIN BUILTIN      BUILTIN      SPU Flow
Fan Tray      Enhanced Fan Tray

node1:
-----
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN124FECOAGB  SRX5600
Midplane      REV 01   760-063936  ACRE2946      Enhanced SRX5600 Midplane
FPM Board     test    710-017254  test          Front Panel Display
PEM 0         Rev 01   740-038514  QCS114111003 DC 2.6kW Power Entry
Module
PEM 1         Rev 01   740-038514  QCS12031100J DC 2.6kW Power Entry

```

```

Module
Routing Engine 0 REV 01 740-056658 9009186342 SRX5k RE-1800X4
CB 0 REV 01 750-062257 CAEB8178 SRX5k SCB3
FPC 0 REV 07 750-044175 CAAD0769 SRX5k SPC II
  CPU BUILTIN BUILTIN SRX5k DPC PPC
  PIC 0 BUILTIN BUILTIN SPU Cp
  PIC 1 BUILTIN BUILTIN SPU Flow
  PIC 2 BUILTIN BUILTIN SPU Flow
  PIC 3 BUILTIN BUILTIN SPU Flow
FPC 4 REV 11 750-043157 CACY1592 SRX5k IOC II
  CPU REV 04 711-043360 CACZ8831 SRX5k MPC PMB
  MIC 1 REV 04 750-049488 CACN0239 10x 10GE SFP+
    PIC 2 BUILTIN BUILTIN 10x 10GE SFP+
      Xcvr 7 REV 01 740-031980 ARN23HW SFP+-10G-SR
      Xcvr 8 REV 01 740-031980 ARN2FVW SFP+-10G-SR
      Xcvr 9 REV 01 740-031980 ARN2YVM SFP+-10G-SR
FPC 5 REV 10 750-056758 CADA8736 SRX5k SPC II
  CPU BUILTIN BUILTIN SRX5k DPC PPC
  PIC 0 BUILTIN BUILTIN SPU Flow
  PIC 1 BUILTIN BUILTIN SPU Flow
  PIC 2 BUILTIN BUILTIN SPU Flow
  PIC 3 BUILTIN BUILTIN SPU Flow
Fan Tray Enhanced Fan Tray

```

show chassis hardware (SRX4200)

```
user@host> show chassis hardware
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			DK2816AR0020	SRX4200
Mainboard	REV 01	650-071675	16061032317	SRX4200
Routing Engine 0		BUILTIN	BUILTIN	SRX Routing Engine
FPC 0		BUILTIN	BUILTIN	FEB
PIC 0		BUILTIN	BUILTIN	8x10G-SFP
Xcvr 0	REV 01	740-038153	MOC11511530020	SFP+-10G-CU3M
Xcvr 1	REV 01	740-038153	MOC11511530020	SFP+-10G-CU3M
Xcvr 2	REV 01	740-038153	MOC11511530020	SFP+-10G-CU3M
Xcvr 3	REV 01	740-038153	MOC11511530020	SFP+-10G-CU3M
Xcvr 4	REV 01	740-021308	04DZ06A00364	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	233363A03066	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	AL70SWE	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	ALN0N6C	SFP+-10G-SR
Xcvr 8	REV 01	740-030076	APF16220018NK1	SFP+-10G-CU1M
Power Supply 0	REV 04	740-041741	1GA26241849	JPSU-650W-AC-AFO
Power Supply 1	REV 04	740-041741	1GA26241846	JPSU-650W-AC-AFO
Fan Tray 0				SRX4200 0, Front to Back
Airflow - AFO				
Fan Tray 1				SRX4200 1, Front to Back
Airflow - AFO				
Fan Tray 2				SRX4200 2, Front to Back
Airflow - AFO				
Fan Tray 3				SRX4200 3, Front to Back
Airflow - AFO				

show chassis hardware clei-models

show chassis hardware clei-models

(SRX5600 and SRX5800 devices with SRX5000 line SRX5K-SCBE [SCB2] and SRX5K-RE-1800X4 [RE2])

```
user@host> show chassis hardware clei-models node 1
node1:
```

Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 01	710-024803		SRX5800-BP-A
FPM Board	REV 01	710-024632		SRX5800-CRAFT-A
PEM 0	Rev 04	740-034724		SRX5800-PWR-4100-AC
PEM 1	Rev 05	740-034724		SRX5800-PWR-4100-AC
Routing Engine 0	REV 01	740-056658	COUCATTBAA	SRX5K-RE-1800X4
CB 0	REV 01	750-056587	COUCATSBAA	SRX5K-SCBE
CB 1	REV 01	750-056587	COUCATSBA	SRX5K-SCBE
CB 2	REV 01	750-056587	COUCATSBA	SRX5K-SCBE
FPC 0	REV 18	750-054877	COUCATLBAA	SRX5K-SPC-4-15-320
CPU		BUILTIN		
FPC 1	REV 18	750-054877	COUCATLBAA	SRX5K-SPC-4-15-320
CPU		BUILTIN		
FPC 2	REV 18	750-054877	COUCATLBAA	SRX5K-SPC-4-15-320
CPU		BUILTIN		
FPC 3	REV 11	750-043157	COUIBCWBAA	SRX5K-MPC
MIC 0	REV 05	750-049486	COUIBCYBAA	SRX-MIC-1X100G-CFP
MIC 1	REV 04	750-049488	COUIBCXBAA	SRX-MIC-10XG-SFPP
FPC 4	REV 18	750-054877	COUCATLBAA	SRX5K-SPC-4-15-320
CPU		BUILTIN		
FPC 7	REV 18	750-054877	COUCATLBAA	SRX5K-SPC-4-15-320
CPU		BUILTIN		
FPC 8	REV 11	750-043157	COUIBCWBAA	SRX5K-MPC
MIC 0	REV 05	750-049486	COUIBCYBAA	SRX-MIC-1X100G-CFP
FPC 9	REV 18	750-054877	COUCATLBAA	SRX5K-SPC-4-15-320
CPU		BUILTIN		
FPC 10	REV 18	750-054877	COUCATLBAA	SRX5K-SPC-4-15-320
CPU		BUILTIN		
Fan Tray 0	REV 04	740-035409		SRX5800-HC-FAN
Fan Tray 1	REV 04	740-035409		SRX5800-HC-FAN

show ethernet-switching mac-learning-log (View)

Supported Platforms	SRX Series
Syntax	show ethernet-switching mac-learning-log
Release Information	Command introduced in Junos OS Release 9.5.
Description	Displays the event log of learned MAC addresses.
Required Privilege Level	view
Related Documentation	
Output Fields	Table 24 on page 349 lists the output fields for the show ethernet-switching mac-learning-log command. Output fields are listed in the approximate order in which they appear.

Table 24: show interfaces Output Fields

Field Name	Field Description
Date and Time	Timestamp when the MAC address was added or deleted from the log.
VLAN-IDX	VLAN index. An internal value assigned by Junos OS for each VLAN.
MAC	Learned MAC address.
Deleted Added	MAC address deleted or added to the MAC learning log.
Blocking	The forwarding state of the interface: <ul style="list-style-type: none"> blocked—Traffic is not being forwarded on the interface. unblocked—Traffic is forwarded on the interface.

Sample Output

show ethernet-switching mac-learning-log

```

user@host> show ethernet-switching mac-learning-log
Wed Mar 18 08:07:05 2009
vlan_idx 7 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 9 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 10 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 11 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 12 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 13 mac 00:00:5E:00:53:00 was deleted

```

```

Wed Mar 18 08:07:05 2009
vlan_idx 14 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 15 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 16 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 4 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 6 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 7 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 9 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 10 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 11 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 12 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 13 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 14 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 15 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 16 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 5 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 18 mac 00:00:5E:00:53:AA was learned
Wed Mar 18 08:07:05 2009
vlan_idx 5 mac 00:00:5E:00:53:AB was learned
Wed Mar 18 08:07:05 2009
vlan_idx 6 mac 00:00:5E:00:53:AC was learned
Wed Mar 18 08:07:05 2009
vlan_idx 16 mac 00:00:5E:00:53:AD was learned
Wed Mar 18 08:07:05 2009
vlan_idx 7 mac 00:00:5E:00:53:AE was learned
Wed Mar 18 08:07:05 2009
vlan_idx 8 mac 00:00:5E:00:53:AF was learned
Wed Mar 18 08:07:05 2009
vlan_idx 12 mac 00:00:5E:00:53:AG was learned
[output truncated]
    
```

show ethernet-switching table (View)

Supported Platforms	SRX Series
Syntax	show ethernet-switching table (brief detail extensive) interface <i>interface-name</i>
Release Information	Command introduced in Junos OS Release 9.5.
Description	Displays the Ethernet switching table.
Options	<ul style="list-style-type: none"> • none—(Optional) Display brief information about the Ethernet switching table. • brief detail extensive—(Optional) Display the specified level of output. • interface-name—(Optional) Display the Ethernet switching table for a specific interface.
Required Privilege Level	view
Related Documentation	
Output Fields	Table 25 on page 351 lists the output fields for the show ethernet-switching table command. Output fields are listed in the approximate order in which they appear.

Table 25: show ethernet-switching table Output Fields

Field Name	Field Description
VLAN	The name of a VLAN.
MAC address	The MAC address associated with the VLAN.
Type	The type of MAC address. Values are: <ul style="list-style-type: none"> • static—The MAC address is manually created. • learn—The MAC address is learned dynamically from a packet's source MAC address. • flood—The MAC address is unknown and flooded to all members.
Age	The time remaining before the entry ages out and is removed from the Ethernet switching table.
Interfaces	Interface associated with learned MAC addresses or All-members (flood entry).
Learned	For learned entries, the time which the entry was added to the Ethernet switching table.

Sample Output

show ethernet-switching table

```
user@host> show ethernet-switching table
Ethernet-switching table: 57 entries, 17 learned
VLAN MAC address Type Age Interfaces
```

```

F2 * Flood - All-members
F2 00:00:5E:00:53:AC Learn 0 ge-0/0/44.0
F2 00:00:5E:00:53:AD Static - Router
Linux * Flood - All-members
Linux 00:00:5E:00:53:AE Static - Router
Linux 00:00:5E:00:53:AF Learn 0 ge-0/0/47.0
T1 * Flood - All-members
T1 00:00:5E:00:53:AA Learn 0 ge-0/0/46.0
T1 00:00:5E:00:53:AB Static - Router
T1 00:00:5E:00:53:AC Learn 0 ge-0/0/46.0
T1 00:00:5E:00:53:AD Static - Router
T10 * Flood - All-members
T10 00:00:5E:00:53:AE Static - Router
T10 00:00:5E:00:53:AF Learn 0 ge-0/0/46.0
T10 00:00:5E:00:53:AG Static - Router
T111 * Flood - All-members
T111 00:00:5E:00:53:AH Learn 0 ge-0/0/15.0
T111 00:00:5E:00:53:AI Static - Router
T111 00:00:5E:00:53:AJ Learn 0 ge-0/0/15.0
T2 * Flood - All-members
T2 00:00:5E:00:53:AK Static - Router
T2 00:00:5E:00:53:AL Learn 0 ge-0/0/46.0
T2 00:00:5E:00:53:AM Static - Router
T3 * Flood - All-members
T3 00:00:5E:00:53:AN Static - Router
T3 00:00:5E:00:53:AO Learn 0 ge-0/0/46.0
T3 00:00:5E:00:53:AP Static - Router
T4 * Flood - All-members
T4 00:00:5E:00:53:AQ Static - Router
T4 00:00:5E:00:53:AR Learn 0 ge-0/0/46.0
[output truncated]

```

Sample Output

show ethernet-switching table brief

```

user@host> show ethernet-switching table brief
Ethernet-switching table: 57 entries, 17 learned
VLAN MAC address Type Age Interfaces
F2 * Flood - All-members
F2 00:00:5E:00:53:AC Learn 0 ge-0/0/44.0
F2 00:00:5E:00:53:AE Static - Router
Linux * Flood - All-members
Linux 00:00:5E:00:53:AA Static - Router
Linux 00:00:5E:00:53:AB Learn 0 ge-0/0/47.0
T1 * Flood - All-members
T1 00:00:5E:00:53:AC Learn 0 ge-0/0/46.0
T1 00:00:5E:00:53:AD Static - Router
T1 00:00:5E:00:53:AE Learn 0 ge-0/0/46.0
T1 00:00:5E:00:53:AF Static - Router
T10 * Flood - All-members
T10 00:00:5E:00:53:AG Static - Router
T10 00:00:5E:00:53:AH Learn 0 ge-0/0/46.0
T10 00:00:5E:00:53:AI Static - Router
T111 * Flood - All-members
T111 00:00:5E:00:53:AJ Learn 0 ge-0/0/15.0
T111 00:00:5E:00:53:AK Static - Router
T111 00:00:5E:00:53:AL Learn 0 ge-0/0/15.0
T2 * Flood - All-members
T2 00:00:5E:00:53:AM Static - Router
T2 00:00:5E:00:53:AN Learn 0 ge-0/0/46.0

```

```

T2 00:00:5E:00:53:A0 Static - Router
T3 * Flood - All-members
T3 00:00:5E:00:53:AP Static - Router
T3 00:00:5E:00:53:AQ Learn 0 ge-0/0/46.0
T3 00:00:5E:00:53:AR Static - Router
T4 * Flood - All-members
T4 00:00:5E:00:53:AS Static - Router
T4 00:00:5E:00:53:AT Learn 0 ge-0/0/46.0
[output truncated]

```

Sample Output

show ethernet-switching table detail

```

user@host> show ethernet-switching table detail
Ethernet-switching table: 57 entries, 17 learned
F2, *
Interface(s): ge-0/0/44.0
Type: Flood
F2, 00:00:5E:00:53:AC
Interface(s): ge-0/0/44.0
Type: Learn, Age: 0, Learned: 2:03:09
F2, 00:00:5E:00:53:AA
Interface(s): Router
Type: Static
Linux, *
Interface(s): ge-0/0/47.0
Type: Flood
Linux, 00:00:5E:00:53:AB
Interface(s): Router
Type: Static
Linux, 00:00:5E:00:53:AC
Interface(s): ge-0/0/47.0
Type: Learn, Age: 0, Learned: 2:03:08
T1, *
Interface(s): ge-0/0/46.0
Type: Flood
T1, 00:00:5E:00:53:AD
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
T1, 00:00:5E:00:53:AE
Interface(s): Router
Type: Static
T1, 00:00:5E:00:53:AF
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
T1, 00:00:5E:00:53:AG
Interface(s): Router
Type: Static
T10, *
Interface(s): ge-0/0/46.0
Type: Flood
T10, 00:00:5E:00:53:AH
Interface(s): Router
Type: Static
T10, 00:00:5E:00:53:AI
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:08
T10, 00:00:5E:00:53:AJ
Interface(s): Router
Type: Static

```

```
T111, *
Interface(s): ge-0/0/15.0
Type: Flood
[output truncated]
```

Sample Output

show ethernet-switching table extensive

```
user@host> show ethernet-switching table extensive
Ethernet-switching table: 57 entries, 17 learned
F2, *
Interface(s): ge-0/0/44.0
Type: Flood
F2, 00:00:5E:00:53:AC
Interface(s): ge-0/0/44.0
Type: Learn, Age: 0, Learned: 2:03:09
F2, 00:00:5E:00:53:AA
Interface(s): Router
Type: Static
Linux, *
Interface(s): ge-0/0/47.0
Type: Flood
Linux, 00:00:5E:00:53:AB
Interface(s): Router
Type: Static
Linux, 00:00:5E:00:53:AC
Interface(s): ge-0/0/47.0
Type: Learn, Age: 0, Learned: 2:03:08
T1, *
Interface(s): ge-0/0/46.0
Type: Flood
T1, 00:00:5E:00:53:AD
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
T1, 00:00:5E:00:53:AE
Interface(s): Router
Type: Static
T1, 00:00:5E:00:53:AF
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
T1, 00:00:5E:00:53:AG
Interface(s): Router
Type: Static
T10, *
Interface(s): ge-0/0/46.0
Type: Flood
T10, 00:00:5E:00:53:AH
Interface(s): Router
Type: Static
T10, 00:00:5E:00:53:AI
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:08
T10, 00:00:5E:00:53:AJ
Interface(s): Router
Type: Static
T111, *
Interface(s): ge-0/0/15.0
Type: Flood
[output truncated]
```

Sample Output

show ethernet-switching table interface ge-0/0/1

```
user@host> show ethernet-switching table interface ge-0/0/1
Ethernet-switching table: 1 unicast entries
VLAN    MAC address      Type    Age Interfaces
V1      *                Flood   - All-members
V1      00:00:5E:00:53:AF Learn    0 ge-0/0/1.0
```

show igmp-snooping route (View)

Supported Platforms	SRX1500, SRX300, SRX320, SRX340, SRX550M, vSRX
Syntax	show igmp-snooping route (brief detail ethernet-switching inet vlan)
Release Information	Command introduced in Junos OS Release 9.5.
Description	Display IGMP snooping route information.
Options	<ul style="list-style-type: none"> • none—Display general parameters. • brief detail—(Optional) Display the specified level of output. • ethernet-switching—(Optional) Display Ethernet switching information. • inet—(Optional) Display inet information. • vlan <i>vlan-id</i> <i>vlan-name</i>—(Optional) Display route information for the specified VLAN.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Understanding Interfaces on page 3
Output Fields	Table 26 on page 356 lists the output fields for the show igmp-snooping route command. Output fields are listed in the approximate order in which they appear.

Table 26: show igmp-snooping route Output Fields

Field Name	Field Description
VLAN	Name of the VLAN.
Group	Multicast group address.
Next-hop	ID associated with the next-hop device.

Sample Output

show igmp-snooping route

```
user@host> show igmp-snooping route
VLAN      Group          Next-hop
v11       203.0.113.0, * 533
Interfaces: ge-0/0/13.0, ge-0/0/1.0
v12       203.0.113.1, * 534
Interfaces: ge-0/0/13.0, ge-0/0/0.0
```

show igmp-snooping route vlan v1

```
user@host> show igmp-snooping route vlan v1
Table: 0
VLAN      Group          Next-hop
```

```
v1      203.0.113.2, *      1266
Interfaces: ge-0/0/0.0
v1      203.0.113.3, *      1266
Interfaces: ge-0/0/0.0
v1      203.0.113.4, *      1266
Interfaces: ge-0/0/0.0
v1      203.0.113.5, *      1266
Interfaces: ge-0/0/0.0
v1      203.0.113.6, *      1266
Interfaces: ge-0/0/0.0
v1      203.0.113.6, *      1266
Interfaces: ge-0/0/0.0
```

show interfaces (SRX Series)

Supported Platforms SRX Series, vSRX

Syntax show interfaces {
 <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 modified in Junos OS Release 9.5.

Description Display status information and statistics about interfaces on SRX Series appliance running Junos OS.

On SRX Series appliance, on configuring identical IPs on a single interface, you will not see a warning message; instead, you will see a syslog message.

- Options**
- **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).
-
- **brief | detail | extensive | terse**—(Optional) Display the specified level of output.
 - **controller**—(Optional) Show controller information.
 - **descriptions**—(Optional) Display interface description strings.
 - **destination-class**—(Optional) Show statistics for destination class.
 - **diagnostics**—(Optional) Show interface diagnostics information.
 - **far-end-interval**—(Optional) Show far end interval statistics.
 - **filters**—(Optional) Show interface filters information.
 - **flow-statistics**—(Optional) Show security flow counters and errors.
 - **interval**—(Optional) Show interval statistics.
 - **load-balancing**—(Optional) Show load-balancing status.
 - **mac-database**—(Optional) Show media access control database information.
 - **mc-ae**—(Optional) Show MC-AE configured interface information.
 - **media**—(Optional) Display media information.
 - **policers**—(Optional) Show interface policers information.
 - **queue**—(Optional) Show queue statistics for this interface.
 - **redundancy**—(Optional) Show redundancy status.
 - **routing**—(Optional) Show routing status.
 - **routing-instance**—(Optional) Name of routing instance.
 - **snmp-index**—(Optional) SNMP index of interface.
 - **source-class**—(Optional) Show statistics for source class.
 - **statistics**—(Optional) Display statistics and detailed output.
 - **switch-port**—(Optional) Front end port number (0..15).
 - **transport**—(Optional) Show interface transport information.
 - **zone**—(Optional) Interface's zone.

Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Understanding Interfaces
List of Sample Output	<p> show interfaces Gigabit Ethernet on page 367 show interfaces brief (Gigabit Ethernet) on page 368 show interfaces detail (Gigabit Ethernet) on page 368 show interfaces extensive (Gigabit Ethernet) on page 370 show interfaces terse on page 373 show interfaces controller (Channelized E1 IQ with Logical E1) on page 373 show interfaces controller (Channelized E1 IQ with Logical DSO) on page 373 show interfaces descriptions on page 374 show interfaces destination-class all on page 374 show interfaces diagnostics optics on page 374 show interfaces far-end-interval coc12-5/2/0 on page 375 show interfaces far-end-interval coc1-5/2/1:1 on page 375 show interfaces filters on page 376 show interfaces flow-statistics (Gigabit Ethernet) on page 376 show interfaces interval (Channelized OC12) on page 377 show interfaces interval (E3) on page 377 show interfaces interval (SONET/SDH) on page 378 show interfaces load-balancing on page 378 show interfaces load-balancing detail on page 378 show interfaces mac-database (All MAC Addresses on a Port) on page 379 show interfaces mac-database (All MAC Addresses on a Service) on page 379 show interfaces mac-database mac-address on page 380 show interfaces mc-ae on page 380 show interfaces media (SONET/SDH) on page 380 show interfaces policers on page 381 show interfaces policers interface-name on page 381 show interfaces queue on page 381 show interfaces redundancy on page 382 show interfaces redundancy (Aggregated Ethernet) on page 382 show interfaces redundancy detail on page 383 show interfaces routing brief on page 383 show interfaces routing detail on page 383 show interfaces routing-instance all on page 384 show interfaces snmp-index on page 384 show interfaces source-class all on page 384 show interfaces statistics (Fast Ethernet) on page 385 show interfaces switch-port on page 385 show interfaces transport pm on page 386 show security zones on page 387 </p>
Output Fields	<p>Table 27 on page 361 lists the output fields for the show interfaces command. Output fields are listed in the approximate order in which they appear.</p>

Table 27: 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
CoS queues	Number of CoS queues configured.	detail extensive none
Current address	Configured MAC address.	detail extensive none

Table 27: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None
Output Rate	Output rate in bps and pps.	None
Active alarms and Active defects	<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 27: 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 <code>ignore-l3-incompletes</code> 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	<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, 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 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 Output Fields (*continued*)

Field Name	Field Description	Level of Output
Ingress queues	Total number of ingress queues supported on the specified interface.	extensive
Queue counters and queue number	<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
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 27: 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 should be 0. 	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. 	extensive
Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number. 	extensive

Table 27: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface. <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive
Interface transmit statistics	Status of the interface-transmit-statistics configuration: Enabled or Disabled.	detail extensive
Queue counters (Egress)	CoS queue number and its associated user-configured forwarding class name. <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
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Traffic statistics	Number and rate of bytes and packets received and transmitted on the specified interface set. <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

Table 27: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Local statistics	Number and rate of bytes and packets destined to the device.	extensive
Transit statistics	Number and rate of bytes and packets transiting the switch. NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the Output bytes and Output packets interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.	extensive
Security	Security zones that interface belongs to.	extensive
Flow Input statistics	Statistics on packets received by flow module.	extensive
Flow Output statistics	Statistics on packets sent by flow module.	extensive
Flow error statistics (Packets dropped due to)	Statistics on errors in the flow module.	extensive
Protocol	Protocol family.	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive none
Flags	Information about protocol family flags. .	detail extensive
Addresses, Flags	Information about the address flags..	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Sample Output

show interfaces 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: 00:1f:12:e4:b1:01, Hardware address: 00:1f:12:e4:b1:01
  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: Sendbcst-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

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
  
```

Sample Output

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: 00:1f:12:e4:b1:01, Hardware address: 00:1f:12:e4:b1:01
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: Sendbcst-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

```

Sample Output

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: 00:1f:12:e4:b1:01, Hardware address: 00:1f:12:e4:b1:01
  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:
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

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
0 best-effort        95 %             950000000      95 %             0 usec          low
none
3 network-control    5 %             50000000       5 %             0 usec          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: Sendbcst-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

```

Sample Output

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
d10                 up   up
d10.0               up   up   inet
dsc                 up   up
gre                 up   up
pip                 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

```

Sample Output

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 DSO)

```

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

```

Sample Output

show interfaces descriptions

```

user@host> show interfaces descriptions
Interface      Admin Link Description
so-1/0/0       up    up    M20-3#1
so-2/0/0       up    up    GSR-12#1
ge-3/0/0       up    up    SMB-OSPF_Area300
so-3/3/0       up    up    GSR-13#1
so-3/3/1       up    up    GSR-13#2
ge-4/0/0       up    up    T320-7#1
ge-5/0/0       up    up    T320-7#2
so-7/1/0       up    up    M160-6#1
ge-8/0/0       up    up    T320-7#3
ge-9/0/0       up    up    T320-7#4
so-10/0/0      up    up    M160-6#2
so-13/0/0      up    up    M20-3#2
so-14/0/0      up    up    GSR-12#2
ge-15/0/0      up    up    SMB-OSPF_Area100
ge-15/0/1      up    up    GSR-13#3
    
```

Sample Output

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)
    
```

Sample Output

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

```

Sample Output

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:

```

Sample Output

show interfaces filters

```

user@host> show interfaces filters
Interface      Admin Link Proto Input Filter      Output Filter
ge-0/0/0       up    up    inet
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
....

```

Sample Output

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
  Isping
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

```

Sample Output

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)

```

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:
  ....

```

Sample Output

show interfaces load-balancing

```

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

```

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

```

Sample Output

show interfaces mac-database (All MAC Addresses on a Port)

```

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)

```

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 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
    
```

Sample Output

show interfaces mc-ae

```

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: 0C48,
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

```

Sample Output

show interfaces policers

```

user@host> show interfaces policers
Interface      Admin Link Proto Input Policer      Output Policer
ge-0/0/0       up    up    inet                ge-0/0/0
ge-0/0/0.0     up    up    inet iso            ge-0/0/0.0
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
so-2/0/0.0     up    up    iso
so-2/1/0       up    down
...

```

show interfaces policers interface-name

```

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
so-2/1/0.0     up    down iso
so-2/1/0.0     up    down inet6

```

Sample Output

show interfaces queue

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
..
..

```

Sample Output

show interfaces redundancy

```

user@host> show interfaces redundancy
Interface State      Last change Primary Secondary Current status
rsp0      Not present
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
r1sq0     On primary   00:06:24 1sq-0/3/0 1sq-1/0/0 both up

```

show interfaces redundancy (Aggregated Ethernet)

```

user@host> show interfaces redundancy
Interface State      Last change Primary Secondary Current status
r1sq0     On secondary 00:56:12 1sq-4/0/0 1sq-3/0/0 both up

ae0
ae1

```

```
ae2
ae3
ae4
```

show interfaces redundancy detail

```
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
```

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
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

```
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
...

```

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   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

```

Sample Output

show interfaces snmp-index

```

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

```

Sample Output

show interfaces source-class all

```

user@host> show interfaces source-class all
Logical interface so-0/1/0.0

          Source class          Packets          Bytes
                                (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
      Source class                Packets                Bytes
                                (packet-per-second)    (bits-per-second)
                                gold                   0                   0
                                (          0) (          0)
                                bronze                  0                   0
                                (          0) (          0)
                                silver                 116113              9753492
                                (          939) (          631616)

```

Sample Output

show interfaces statistics (Fast Ethernet)

```

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

```

Sample Output

show interfaces switch-port

```

user@host# show interfaces ge-slot/0/0 switch-port port-number
Port 0, Physical link is Up
  Speed: 100mbps, Auto-negotiation: Enabled
  Statistics:
    Total bytes                Receive                Transmit
    28437086                    28437086              21792250

```

```

Total packets          409145          88008
Unicast packets        9987          83817
Multicast packets      145002         0
Broadcast packets     254156         4191
Multiple collisions    23            10
FIFO/CRC/Align errors 0              0
MAC pause frames       0              0
Oversized frames       0
Runt frames            0
Jabber frames          0
Fragment frames        0
Discarded frames       0
Autonegotiation information:
Negotiation status: Complete
Link partner:
Link mode: Full-duplex, Flow control: None, Remote fault: OK, Link
partner Speed: 100 Mbps
Local resolution:
Flow control: None, Remote fault: Link OK

```

Sample Output

show interfaces transport pm

```

user@host> show interfaces transport pm all current et-0/1/0
Physical interface: et-0/1/0, SNMP ifIndex 515
14:45-current Elapse time:900 Seconds
Near End      Suspect Flag:False      Reason:None
PM            COUNT          THRESHOLD    TCA-ENABLED    TCA-RAISED

OTU-BBE       0              800          No              No
OTU-ES        0              135          No              No
OTU-SES       0              90           No              No
OTU-UAS       427           90           No              No
Far End      Suspect Flag:True      Reason:Unknown
PM            COUNT          THRESHOLD    TCA-ENABLED    TCA-RAISED

OTU-BBE       0              800          No              No
OTU-ES        0              135          No              No
OTU-SES       0              90           No              No
OTU-UAS       0              90           No              No
Near End      Suspect Flag:False      Reason:None
PM            COUNT          THRESHOLD    TCA-ENABLED    TCA-RAISED

ODU-BBE       0              800          No              No
ODU-ES        0              135          No              No
ODU-SES       0              90           No              No
ODU-UAS       427           90           No              No
Far End      Suspect Flag:True      Reason:Unknown
PM            COUNT          THRESHOLD    TCA-ENABLED    TCA-RAISED

ODU-BBE       0              800          No              No
ODU-ES        0              135          No              No
ODU-SES       0              90           No              No
ODU-UAS       0              90           No              No
FEC           Suspect Flag:False      Reason:None
PM            COUNT          THRESHOLD    TCA-ENABLED    TCA-RAISED

FEC-CorrectedErr 2008544300    0            NA              NA
FEC-UncorrectedWords 0              0            NA              NA
BER           Suspect Flag:False      Reason:None

```

```

PM          MIN      MAX      AVG      THRESHOLD      TCA-ENABLED
TCA-RAISED
BER          3.6e-5   5.8e-5   3.6e-5   10.0e-3         No
Yes
Physical interface: et-0/1/0, SNMP ifIndex 515
14:45-current
Suspect Flag: True          Reason: Object Disabled
PM          CURRENT  MIN      MAX      AVG      THRESHOLD
TCA-ENABLED      TCA-RAISED
(MIN)
(MAX) (MIN) (MAX) (MIN) (MAX)
Lane chromatic dispersion          0      0      0      0
0      NA      NA      NA      NA
Lane differential group delay      0      0      0      0
0      NA      NA      NA      NA
q Value          120      120      120      120      0
0      NA      NA      NA      NA
SNR          28      28      29      28      0
0      NA      NA      NA      NA
Tx output power(0.01dBm)          -5000 -5000 -5000 -5000 -300
-100      No      No      No      No
Rx input power(0.01dBm)          -3642 -3665 -3626 -3637 -1800
-500      No      No      No      No
Module temperature(Celsius)      46      46      46      46      -5
75      No      No      No      No
Tx laser bias current(0.1mA)      0      0      0      0      0
0      NA      NA      NA      NA
Rx laser bias current(0.1mA)      1270   1270   1270   1270   0
0      NA      NA      NA      NA
Carrier frequency offset(MHz)      -186   -186   -186   -186   -5000
5000      No      No      No      No

```

Sample Output

show security zones

```

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 diagnostics optics

Supported Platforms [SRX Series, vSRX](#)

Syntax show interfaces diagnostics optics *interface-name*

Release Information Command introduced in Junos OS Release 10.1.

Description Display diagnostics data and alarms for Gigabit Ethernet optical transceivers (SFP) installed in SRX Series Services Gateways. The information provided by this command is known as digital optical monitoring (DOM) information.

Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transponder vendors. Generally, a high alarm or low alarm indicates that the optics module is not operating properly. This information can be used to diagnose why a transceiver is not working.

Options *interface-name*—Name of the interface associated with the port in which the transceiver is installed: *ge-fpc/pic/port* .

Required Privilege Level view

Related Documentation

- [Understanding Interfaces on page 3](#)

List of Sample Output [show interfaces diagnostics optics on page 392](#)

Output Fields [Table 28 on page 389](#) lists the output fields for the show interfaces diagnostics optics command. Output fields are listed in the general order in which they appear.

Table 28: show interfaces diagnostics optics Output Fields

Field Name	Field Description
Physical interface	Displays the name of the physical interface.
Laser bias current	Displays the magnitude of the laser bias power setting current, in milliamperes. The laser bias provides direct modulation of laser diodes and modulates currents.
Laser output power	Displays the laser output power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Module temperature	Displays the temperature, in Celsius and Fahrenheit.
Module voltage	Displays the voltage, in Volts.
Receiver signal average optical power	Displays the receiver signal average optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).

Table 28: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Laser bias current high alarm	Displays whether the laser bias power setting high alarm is On or Off .
Laser bias current low alarm	Displays whether the laser bias power setting low alarm is On or Off .
Laser bias current high warning	Displays whether the laser bias power setting high warning is On or Off .
Laser bias current low warning	Displays whether the laser bias power setting low warning is On or Off .
Laser output power high alarm	Displays whether the laser output power high alarm is On or Off .
Laser output power low alarm	Displays whether the laser output power low alarm is On or Off .
Laser output power high warning	Displays whether the laser output power high warning is On or Off .
Laser output power low warning	Displays whether the laser output power low warning is On or Off .
Module temperature high alarm	Displays whether the module temperature high alarm is On or Off .
Module temperature low alarm	Displays whether the module temperature low alarm is On or Off .
Module temperature high warning	Displays whether the module temperature high warning is On or Off .
Module temperature low warning	Displays whether the module temperature low warning is On or Off .
Module voltage high alarm	Displays whether the module voltage high alarm is On or Off .
Module voltage low alarm	Displays whether the module voltage low alarm is On or Off .
Module voltage high warning	Displays whether the module voltage high warning is On or Off .
Module voltage low warning	Displays whether the module voltage low warning is On or Off .
Laser rx power high alarm	Displays whether the receive laser power high alarm is On or Off .

Table 28: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Laser rx power low alarm	Displays whether the receive laser power low alarm is On or Off .
Laser rx power high warning	Displays whether the receive laser power high warning is On or Off .
Laser rx power low warning	Displays whether the receive laser power low warning is On or Off .
Laser bias current high alarm threshold	Displays the vendor-specified threshold for the laser bias current high alarm.
Laser bias current low alarm threshold	Displays the vendor-specified threshold for the laser bias current low alarm.
Laser bias current high warning threshold	Displays the vendor-specified threshold for the laser bias current high warning.
Laser bias current low warning threshold	Displays the vendor-specified threshold for the laser bias current low warning.
Laser output power high alarm threshold	Displays the vendor-specified threshold for the laser output power high alarm.
Laser output power low alarm threshold	Displays the vendor-specified threshold for the laser output power low alarm.
Laser output power high warning threshold	Displays the vendor-specified threshold for the laser output power high warning.
Laser output power low warning threshold	Displays the vendor-specified threshold for the laser output power low warning.
Module temperature high alarm threshold	Displays the vendor-specified threshold for the module temperature high alarm.
Module temperature low alarm threshold	Displays the vendor-specified threshold for the module temperature low alarm.
Module temperature high warning threshold	Displays the vendor-specified threshold for the module temperature high warning.
Module temperature low warning threshold	Displays the vendor-specified threshold for the module temperature low warning.
Module voltage high alarm threshold	Displays the vendor-specified threshold for the module voltage high alarm.
Module voltage low alarm threshold	Displays the vendor-specified threshold for the module voltage low alarm.

Table 28: show interfaces diagnostics optics Output Fields (continued)

Field Name	Field Description
Module voltage high warning threshold	Displays the vendor-specified threshold for the module voltage high warning.
Module voltage low warning threshold	Displays the vendor-specified threshold for the module voltage low warning.
Laser rx power high alarm threshold	Displays the vendor-specified threshold for the laser rx power high alarm.
Laser rx power low alarm threshold	Displays the vendor-specified threshold for the laser rx power low alarm.
Laser rx power high warning threshold	Displays the vendor-specified threshold for the laser rx power high warning.
Laser rx power low warning threshold	Displays the vendor-specified threshold for the laser rx power low warning.

Sample Output

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 flow-statistics

Supported Platforms [SRX Series, vSRX](#)

Syntax `show interfaces flow-statistics <interface-name>`

Release Information Command introduced in Junos OS Release 9.2.

Description Display interfaces flow statistics.

Options *Interface-name* —(Optional) Display flow statistics about the specified interface. Following is a list of typical interface names. Replace *pim* with the PIM slot and *port* with the port number. For a complete list, see the [“Interface Naming Conventions”](#) on page 8.

- *at-pim/O/port*—ATM-over-ADSL or ATM-over-SHDSL interface.
- *br-pim/O/port*—Basic Rate Interface for establishing ISDN connections.
- *ce1-pim/O/port*—Channelized E1 interface.
- *ct1-pim/O/port*—Channelized T1 interface.
- *dl0*—Dialer Interface for initiating ISDN and USB modem connections.
- *e1-pim/O/port*—E1 interface.
- *e3-pim/O/port*—E3 interface.
- *fe-pim/O/port*—Fast Ethernet interface.
- *ge-pim/O/port*—Gigabit Ethernet interface.
- *se-pim/O/port*—Serial interface.
- *t1-pim/O/port*—T1 (also called DS1) interface.
- *t3-pim/O/port*—T3 (also called DS3) interface.
- *wx-slot/O/O*—WAN acceleration interface, for the WXC Integrated Services Module (ISM 200).

Required Privilege Level view

Related Documentation

- [Juniper Networks Devices Processing Overview](#)
- [Understanding Interfaces](#) on page 3

List of Sample Output [show interfaces flow-statistics \(Gigabit Ethernet\)](#) on page 397

Output Fields [Table 29](#) on page 395 lists the output fields for the `show interfaces flow-statistics` command. Output fields are listed in the approximate order in which they appear.

Table 29: show interfaces flow-statistics Output Fields

Field Name	Field Description
Traffic statistics	Number of packets and bytes transmitted and received on the physical interface.
Local statistics	Number of packets and bytes transmitted and received on the physical interface.
Transit statistics	Number of packets and bytes transiting the physical interface.
Flow input statistics	Statistics on packets received by flow module.
Flow output statistics	Statistics on packets sent by flow module.
Flow error statistics	Packet drop statistics for the flow module. For further details, see Table 30 on page 395 .

Table 30: Flow Error Statistics (Packet Drop Statistics for the Flow Module)

Error	Error Description
Screen:	
Address spoofing	The packet was dropped when the screen module detected address spoofing.
Syn-attack protection	The packet was dropped because of SYN attack protection or SYN cookie protection.
VPN:	
Authentication failed	The packet was dropped because the IPsec Encapsulating Security Payload (ESP) or Authentication Header (AH) authentication failed.
No SA for incoming SPI	The packet was dropped because the incoming IPsec packet's security parameter index (SPI) does not match any known SPI.
Security association not active	The packet was dropped because an IPsec packet was received for an inactive SA.
NAT:	
Incoming NAT errors	The source NAT rule search failed, an invalid source NAT binding was found, or the NAT allocation failed.
Multiple incoming NAT	Sometimes packets are looped through the system more than once; if source NAT is specified more than once, the packet will be dropped.
Auth:	
Multiple user authentications	Sometimes packets are looped through the system more than once. Each time a packet passes through the system, that packet must be permitted by a policy. If the packet matches more than one policy that specifies user authentication, then it will be dropped.

Table 30: Flow Error Statistics (Packet Drop Statistics for the Flow Module) (*continued*)

User authentication errors	<p>Packet was dropped because policy requires authentication; however:</p> <ul style="list-style-type: none"> • Only Telnet, FTP, and HTTP traffic can be authenticated. • The corresponding authentication entry could not be found, if web-auth is specified. • The maximum number of authenticated sessions per user was exceeded.
Flow:	
No one interested in self packets	<p>This counter is incremented for one of the following reasons:</p> <ul style="list-style-type: none"> • The outbound interface is a self interface, but the packet is not marked as a to-self packet and the destination address is in a source NAT pool. • No service is interested in the to-self packet • When a zone has ident-reset service enabled, the TCP RST to IDENT request for port 113 is sent back and this counter is incremented.
No minor session	The packet was dropped because no minor sessions are available and a minor session was requested. Minor sessions are allocated for storing additional TCP state information.
No more sessions	The packet was dropped because there were no more free sessions available.
No route present	<p>The packet was dropped because a valid route was not available to forward the packet.</p> <p>For new sessions, the counter is incremented for one of the following reasons:</p> <ul style="list-style-type: none"> • No valid route was found to forward the packet. • A discard or reject route was found. • The route could not be added due to lack of memory. • The reverse path forwarding check failed for an incoming multicast packet. <p>For existing sessions, the prior route was changed or deleted, or a more specific route was added. The session is rerouted, and this reroute could fail because:</p> <ul style="list-style-type: none"> • A new route could not be found; either the previous route was removed, or the route was changed to discard or reject. • Multiple packets may concurrently force rerouting to occur, and only one packet can successfully complete the rerouting process. Other packets will be dropped. • The route table was locked for updates by the Routing Engine. Packets that match a new session are retried, whereas packets that match an existing session are not.
No tunnel found	The packet was dropped because a valid tunnel could not be found
No session for a gate	This counter is incremented when a packet is destined for an ALG, and the ALG decides to drop this packet.
No zone or NULL zone binding	The packet was dropped because its incoming interface was not bound to any zone.
Policy denied	<p>The error counter is incremented for one of the following reasons:</p> <ul style="list-style-type: none"> • Source and/or destination NAT has occurred and policy says to drop the packet. • Policy specifies user authentication, which failed. • Policy was configured to deny this packet.

Table 30: Flow Error Statistics (Packet Drop Statistics for the Flow Module) (*continued*)

TCP sequence number out of window	A TCP packet with a sequence number failed the TCP sequence number check that was received.
Counters Not Currently in Use	
No parent for a gate	-
Invalid zone received packet	-
No NAT gate	-

Sample Output

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 igmp 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 queue

Supported Platforms vSRX

Syntax show interfaces queue
 <both-ingress-egress>
 <egress>
 <forwarding-class *forwarding-class*>
 <ingress>
 <interface-name *interface-name*>
 <l2-statistics>

Release Information Command introduced in Junos OS Release 15.1X49-D30 for vSRX.

Description Display class-of-service (CoS) queue information for physical interfaces.

Options none—Show detailed CoS queue statistics for all physical interfaces.

both-ingress-egress—Display both ingress and egress queue statistics.

egress—Display egress queue statistics.

forwarding-class *forwarding-class*—(Optional) Forwarding class name for this queue. Show detailed CoS statistics for the queue that is associated with the specified forwarding class.

ingress—Display ingress queue statistics.

interface-name *interface-name*—(Optional) Show detailed CoS queue statistics for the specified interface.

l2-statistics—(Optional) Display Layer 2 statistics for MLPPP, FRF.15, and FRF.16 bundles.

Required Privilege Level view

Related Documentation

- [Understanding Class of Service](#)

List of Sample Output [show interfaces queue \(vSRX\) on page 401](#)

Output Fields [Table 31 on page 399](#) lists the output fields for the **show interfaces queue** command. Output fields are listed in the approximate order in which they appear.

Table 31: show interfaces queue Output Fields

Field Name	Field Description
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .

Table 31: show interfaces queue Output Fields (*continued*)

Field Name	Field Description
Interface index	Index number of the physical interface. The number reflects the interface's initialization sequence.
SNMP ifIndex	SNMP index number for the interface.
Forwarding classes supported	Total number of forwarding classes supported on the specified interface.
Forwarding classes in use	Total number of forwarding classes in use on the specified interface.
Egress queues supported	Total number of egress queues supported on the specified interface.
Egress queues in use	Total number of egress queues in use on the specified interface.
The following output fields are applicable to both the interface component and Packet Forwarding Engine component in the <code>show interfaces queue</code> command:	
Queue	Queue number.
Forwarding classes	Forwarding class name.
Queued Packets	Number of packets in this queue.
Queued Bytes	Number of bytes in this queue.
Transmitted Packets	Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the Packet Forwarding Engine Chassis Queues field) shows the prefragmentation values.
Transmitted Bytes	Number of bytes transmitted by this queue.
Tail-dropped packets	Number of packets dropped because of tail drop.
RL-dropped bytes	Number of bytes dropped because of rate limiting.
RED-dropped packets	Number of packets dropped because of random early detection (RED).
RED-dropped bytes	Number of bytes dropped because of RED. <ul style="list-style-type: none"> • Low, non-TCP—Number of low-loss priority, non-TCP bytes dropped because of RED. • Low, TCP—Number of low-loss priority, TCP bytes dropped because of RED. • High, non-TCP—Number of high-loss priority, non-TCP bytes dropped because of RED. • High, TCP—Number of high-loss priority, TCP bytes dropped because of RED.
Queue Buffer Usage:	<ul style="list-style-type: none"> • Reserved buffer—The size of the memory buffer that is allocated for storing packets • Current—The amount of buffer memory that is currently in use on this queue.

Sample Output

show interfaces queue (vSRX)

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 statistics (View)

Supported Platforms [SRX Series, vSRX](#)

Syntax `show interfaces statistics interface-name`

Release Information Command introduced in Junos OS Release 10.1.

Description Displays the interface input and output statistics for physical and logical interface.

Required Privilege Level view

Related Documentation

- [Understanding Interfaces on page 3](#)

List of Sample Output [show interfaces statistics on page 402](#)

Sample Output

show interfaces statistics

```
user@host> show interfaces statistics st0.1
Logical interface st0.1 (Index 91) (SNMP ifIndex 268)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Secure-Tunnel
  Input packets : 2743333
  Output packets: 6790470992
  Security: Zone: untrust
  Allowed host-inbound traffic : bootp bfd bgp dns dvmrp igmp ldp msdp nhrp
  ospf pgm pim rip router-discovery rsvp sap vrrp dhcp finger ftp tftp ident-reset
  http https ike netconf ping reverse-telnet
  reverse-ssh rlogin rpm rsh snmp snmp-trap ssh telnet traceroute xnm-clear-text
  xnm-ssl lsping ntp sip
  Protocol inet, MTU: 9192
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 192.167.1.0/30, Local: 192.167.1.1
```

show interfaces terse zone

Supported Platforms [SRX Series](#)

Syntax show interfaces terse zone

Release Information Command introduced in Junos OS Release 12.3X48-D20.

Description Display summary information about zone interfaces.

Options This command has no options.

Required Privilege Level view

Sample Output

show interface terse zone

```
user@host> show interface terse zone
Interface      Admin  Link  Proto  Local          Remote      Zone
ge-0/0/0.0    up     up    inet   1.4.253.251/16  trust
```

show ipv6 neighbors

Supported Platforms	SRX1500, SRX320, SRX340, SRX550M, vSRX
Syntax	show ipv6 neighbors
Release Information	Command introduced in Junos OS Release 12.1X45-D10.
Description	Display information about the IPv6 neighbor cache.
Options	This command has no options.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear ipv6 neighbors on page 323
List of Sample Output	show ipv6 neighbors on page 404
Output Fields	Table 32 on page 404 lists the output fields for the show ipv6 neighbors command. Output fields are listed in the approximate order in which they appear.

Table 32: show ipv6 neighbors Output Fields

Field Name	Field Description
IPv6 Address	Name of the IPv6 interface.
Linklayer Address	Link-layer address.
State	State of the link: up, down, incomplete, reachable, stale, or unreachable.
Exp	Number of seconds until the entry expires.
Rtr	Whether the neighbor is a routing device: yes or no.
Secure	Whether this entry was created using the Secure Neighbor Discovery (SEND) protocol: yes or no.
Interface	Name of the interface.

Sample Output

show ipv6 neighbors

```

user@host> show ipv6 neighbors
IPv6 Address      Linklayer Address  State      Exp Rtr Secure Interface
10:1::2          00:00:0a:00:00:00  reachable  17  yes no   reth0.0
11:11::2         00:19:e2:4b:61:83  stale      1197 yes no   at-1/0/0.0

```

```
12:12::2      00:19:e2:4b:61:83  stale  1188 yes no  at-3/0/0.0
```

show lacp interfaces (View)

Supported Platforms [SRX Series](#)

Syntax `show lacp interfaces interface-name`

Release Information Command modified in Junos OS Release 10.2.

Description Display Link Aggregation Control Protocol (LACP) information about the specified aggregated Ethernet interface, redundant Ethernet interface, Gigabit Ethernet interface, or 10-Gigabit Ethernet interface. If you do not specify an interface name, LACP information for all interfaces is displayed.

Options `none`—Display LACP information for all interfaces.

interface-name—(Optional) Display LACP information for the specified interface:

- Aggregated Ethernet—*aenumber*
- Redundant Ethernet—*rethnumber*
- Gigabit Ethernet—*ge-fpc/pic/port*
- 10-Gigabit Ethernet—*xe-fpc/pic/port*



NOTE: The `show lacp interfaces` command returns the following error message if your system is not configured in either active or passive LACP mode:

“Warning: lacp subsystem not running – not needed by configuration”

Required Privilege Level view

Related Documentation [• Verifying LACP on Redundant Ethernet Interfaces on page 157](#)

List of Sample Output [show lacp interfaces \(Aggregated Ethernet\) on page 408](#)
[show lacp interfaces \(Redundant Ethernet\) on page 409](#)
[show lacp interfaces \(Gigabit Ethernet\) on page 409](#)

Output Fields [Table 33 on page 406](#) lists the output fields for the `show lacp interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 33: show lacp interfaces Output Fields

Field Name	Field Description
Aggregated interface	Aggregated interface value.

Table 33: show lacp interfaces Output Fields (*continued*)

Field Name	Field Description
LACP State	<p>LACP state information for each aggregated interface:</p> <ul style="list-style-type: none"> • Role—Role played by the interface. It can be one of the following: <ul style="list-style-type: none"> • Actor—Local device participating in LACP negotiation. • Partner—Remote device participating in LACP negotiation. • Exp—Expired state. Yes indicates the actor or partner is in an expired state. No indicates the actor or partner is not in an expired state. • Def—Default. Yes indicates that the actor's receive machine is using the default operational partner information, administratively configured for the partner. No indicates the operational partner information in use has been received in a link aggregation control protocol data unit (PDU). • Dist—Distribution of outgoing frames. No indicates distribution of outgoing frames on the link is currently disabled and is not expected to be enabled. Otherwise, the value is Yes. • Col—Collection of incoming frames. Yes indicates collection of incoming frames on the link is currently enabled and is not expected to be disabled. Otherwise, the value is No. • Syn—Synchronization. If the value is Yes, the link is considered synchronized. It has been allocated to the correct link aggregation group, the group has been associated with a compatible aggregator, and the identity of the link aggregation group is consistent with the system ID and operational key information transmitted. If the value is No, the link is not synchronized. It is currently not in the right aggregation. • Aggr—Ability of aggregation port to aggregate (Yes) or to operate only as an individual link (No). • Timeout—LACP timeout preference. Periodic transmissions of link aggregation control PDUs occur at either a slow or fast transmission rate, depending upon the expressed LACP timeout preference (Long Timeout or Short Timeout). • Activity—Actor or partner's port activity. Passive indicates the port's preference for not transmitting link aggregation control PDUs unless its partner's control value is Active. Active indicates the port's preference to participate in the protocol regardless of the partner's control value.

Table 33: show lacp interfaces Output Fields (*continued*)

Field Name	Field Description
LACP Protocol	<p>LACP protocol information for each aggregated interface:</p> <ul style="list-style-type: none"> Link state (active or standby) indicated in parentheses next to the interface when link protection is configured. Receive State—One of the following values: <ul style="list-style-type: none"> Current—The state machine receives a link aggregation control PDU and enters the Current state. Defaulted—If no link aggregation control PDU is received before the timer for the Current state expires a second time, the state machine enters the Defaulted state. Expired—If no link aggregation control PDU is received before the timer for the Current state expires once, the state machine enters the Expired state. Initialize—When the physical connectivity of a link changes or a Begin event occurs, the state machine enters the Initialize state. LACP Disabled—If the port is operating in half duplex, the operation of LACP is disabled on the port, forcing the state to LACP Disabled. This state is similar to the Defaulted state, except that the port is forced to operate as an individual port. Port Disabled—If the port becomes inoperable and a Begin event has not occurred, the state machine enters the Port Disabled state. Transmit State—Transmit state of state machine. One of the following values: <ul style="list-style-type: none"> Fast Periodic—Periodic transmissions are enabled at a fast transmission rate. No Periodic—Periodic transmissions are disabled. Periodic Timer—Transitory state entered when the periodic timer expires. Slow Periodic—Periodic transmissions are enabled at a slow transmission rate. Mux State—State of the multiplexer state machine for the aggregation port. The state is one of the following values: <ul style="list-style-type: none"> Attached—Multiplexer state machine initiates the process of attaching the port to the selected aggregator. Collecting Distributing—Collecting and distributing states are merged together to form a combined state (coupled control). Because independent control is not possible, the coupled control state machine does not wait for the partner to signal that collection has started before enabling both collection and distribution. Detached—Process of detaching the port from the aggregator is in progress. Waiting—Multiplexer state machine is in a holding process, awaiting an outcome.

Sample Output

show lacp interfaces (Aggregated Ethernet)

```

user@host> show lacp interfaces ae0
Aggregated interface: ae0
LACP state:      Role   Exp   Def   Dist  Col   Syn   Aggr  Timeout  Activity
ge-2/0/0        Actor  No    No    Yes   Yes   Yes   Yes     Fast     Active
ge-2/0/0        Partner No    No    Yes   Yes   Yes   Yes     Fast     Active
ge-2/0/1        Actor  No    No    Yes   Yes   Yes   Yes     Fast     Active
ge-2/0/1        Partner No    No    Yes   Yes   Yes   Yes     Fast     Active
ge-2/2/0        Actor  No    No    Yes   Yes   Yes   Yes     Fast     Active
ge-2/2/0        Partner No    No    Yes   Yes   Yes   Yes     Fast     Active
ge-2/2/1        Actor  No    No    Yes   Yes   Yes   Yes     Fast     Active
ge-2/2/1        Partner No    No    Yes   Yes   Yes   Yes     Fast     Active
LACP protocol:  Receive State  Transmit State  Mux State
ge-2/0/0        Current      Fast periodic  Collecting distributing

```

```

ge-2/0/1          Current  Fast periodic Collecting distributing
ge-2/2/0          Current  Fast periodic Collecting distributing
ge-2/2/1          Current  Fast periodic Collecting distributing

```

show lacp interfaces (Redundant Ethernet)

```

user@host> show lacp interfaces reth0
Aggregated interface: reth0
LACP state:
  Role  Exp  Def  Dist  Co1  Syn  Aggr  Timeout  Activity
ge-11/0/0 Actor No  No  Yes  Yes  Yes  Yes  Fast  Active
ge-11/0/0 Partner No  No  Yes  Yes  Yes  Yes  Fast  Active
ge-11/0/1 Actor No  No  Yes  Yes  Yes  Yes  Fast  Active
ge-11/0/1 Partner No  No  Yes  Yes  Yes  Yes  Fast  Active
ge-11/0/2 Actor No  No  Yes  Yes  Yes  Yes  Fast  Active
ge-11/0/2 Partner No  No  Yes  Yes  Yes  Yes  Fast  Active
ge-11/0/3 Actor No  No  Yes  Yes  Yes  Yes  Fast  Active
ge-11/0/3 Partner No  No  Yes  Yes  Yes  Yes  Fast  Active
ge-3/0/0 Actor No  No  Yes  Yes  Yes  Yes  Fast  Active
ge-3/0/0 Partner No  No  Yes  Yes  Yes  Yes  Fast  Active
ge-3/0/1 Actor No  No  Yes  Yes  Yes  Yes  Fast  Active
ge-3/0/1 Partner No  No  Yes  Yes  Yes  Yes  Fast  Active
ge-3/0/2 Actor No  No  Yes  Yes  Yes  Yes  Fast  Active
ge-3/0/2 Partner No  No  Yes  Yes  Yes  Yes  Fast  Active
ge-3/0/3 Actor No  No  Yes  Yes  Yes  Yes  Fast  Active
ge-3/0/3 Partner No  No  Yes  Yes  Yes  Yes  Fast  Active
LACP protocol:
  Receive State  Transmit State  Mux State
ge-11/0/0        Current        Fast periodic Collecting distributing
ge-11/0/1        Current        Fast periodic Collecting distributing
ge-11/0/2        Current        Fast periodic Collecting distributing
ge-11/0/3        Current        Fast periodic Collecting distributing
ge-3/0/0         Current        Fast periodic Collecting distributing
ge-3/0/1         Current        Fast periodic Collecting distributing
ge-3/0/2         Current        Fast periodic Collecting distributing
ge-3/0/3         Current        Fast periodic Collecting distributing
{primary:node1}

```

show lacp interfaces (Gigabit Ethernet)

```

user@host> show lacp interfaces ge-0/3/0
Aggregated interface: ae0
LACP State:
  Role  Exp  Def  Dist  Co1  Syn  Aggr  Timeout  Activity
ge-0/3/0 Actor No  No  Yes  Yes  Yes  Yes  Fast  Active
ge-0/3/0 Partner No  No  Yes  Yes  Yes  Yes  Fast  Active
LACP Protocol:
  Receive State  Transmit State  Mux State
ge-0/3/0        Current        Fast periodic Collecting distributing

```

show lacp statistics interfaces (View)

Supported Platforms [SRX Series](#)

Syntax `show lacp statistics interfaces interface-name`

Release Information Command modified in Junos OS Release 10.2.

Description Display Link Aggregation Control Protocol (LACP) statistics about the specified aggregated Ethernet interface or redundant Ethernet interface. If you do not specify an interface name, LACP statistics for all interfaces are displayed.

Options *interface-name*—(Optional) Name of an interface.

Required Privilege Level view

Related Documentation

- [Verifying LACP on Redundant Ethernet Interfaces on page 157](#)

List of Sample Output [show lacp statistics interfaces on page 410](#)

Output Fields [Table 34 on page 410](#) lists the output fields for the `show lacp statistics interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 34: show lacp statistics interfaces Output Fields

Field Name	Field Description
Aggregated interface	Aggregated interface value.
LACP Statistics	<p>LACP statistics provide the following information:</p> <ul style="list-style-type: none"> • LACP Rx—LACP received counter that increments for each normal hello. • LACP Tx—Number of LACP transmit packet errors logged. • Unknown Rx—Number of unrecognized packet errors logged. • Illegal Rx—Number of invalid packets received.

Sample Output

show lacp statistics interfaces

```

user@host> show lacp statistics interfaces ae0
Aggregated interface: ae0
LACP Statistics:      LACP Rx      LACP Tx      Unknown Rx      Illegal Rx
ge-2/0/0              1352         2035          0                0
ge-2/0/1              1352         2056          0                0
ge-2/2/0              1352         2045          0                0
ge-2/2/1              1352         2043          0                0

```

show oam ethernet link-fault-management

Supported Platforms	SRX1500, SRX300, SRX320, SRX340, SRX550M
Syntax	show oam ethernet link-fault-management <brief detail> <interface-name>
Release Information	Statement for branch SRX Series devices introduced in Junos OS Release 9.5.
Description	Display Operation, Administration, and Maintenance (OAM) link fault management (LFM) information for Ethernet interfaces.
Options	brief detail —(Optional) Display the specified level of output. interface-name —(Optional) Display link fault management information for the specified Ethernet interface only.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear oam ethernet connectivity-fault-management path-database on page 319 • clear oam ethernet connectivity-fault-management statistics • Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways on page 181 • Example: Configuring Ethernet OAM Link Fault Management on page 183
List of Sample Output	show oam ethernet link-fault-management brief on page 415 show oam ethernet link-fault-management detail on page 415
Output Fields	Table 35 on page 411 lists the output fields for the show oam ethernet link-fault-management command. Output fields are listed in the approximate order in which they appear.

Table 35: show oam ethernet link-fault-management Output Fields

Field Name	Field Description	Level of Output
Status	Status of the established link. <ul style="list-style-type: none"> • Fail—A link fault condition exists. • Running—A link fault condition does not exist. 	All levels
Discovery state	State of the discovery mechanism: <ul style="list-style-type: none"> • Passive Wait • Send Any • Send Local Remote • Send Local Remote Ok 	All levels

Table 35: show oam ethernet link-fault-management Output Fields (*continued*)

Field Name	Field Description	Level of Output
Peer address	Address of the OAM peer.	All levels
Flags	<p>Information about the interface.</p> <ul style="list-style-type: none"> • Remote-Stable—Indicates remote OAM client acknowledgment of, and satisfaction with, local OAM state information. False indicates that remote DTE has either not seen or is unsatisfied with local state information. True indicates that remote DTE has seen and is satisfied with local state information. • Local-Stable—Indicates local OAM client acknowledgment of, and satisfaction with, remote OAM state information. False indicates that local DTE either has not seen or is unsatisfied with remote state information. True indicates that local DTE has seen and is satisfied with remote state information. • Remote-State-Valid—Indicates the OAM client has received remote state information found within local information TLVs (type, length, values) of received Information OAM PDUs. False indicates that the OAM client has not seen remote state information. True indicates that the OAM client has seen remote state information. 	All levels
Remote loopback status	An OAM entity can put its remote peer into loopback mode using the Loopback control OAM PDU. In loopback mode, every frame received is transmitted back on the same port (except for OAM PDUs, which are needed to maintain the OAM session).	All levels
Remote entity information	<p>Remote entity information.</p> <ul style="list-style-type: none"> • Remote MUX action—Indicates the state of the multiplexer functions of the OAM sublayer. Device is forwarding non-OAM PDUs to the lower sublayer or discarding non-OAM PDUs. • Remote parser action—Indicates the state of the parser function of the OAM sublayer. Device is forwarding non-OAM PDUs to the higher sublayer, looping back non-OAM PDUs to the lower sublayer, or discarding non-OAM PDUs. • Discovery mode—Indicates whether discovery mode is active or inactive. • Unidirectional mode—Indicates the ability to operate a link in unidirectional mode for diagnostic purposes. • Remote loopback mode—Indicates whether remote loopback is supported or not supported. • Link events—Indicates whether interpreting link events is supported or not supported on the remote peer. • Variable requests—Indicates whether variable requests are supported or not supported. The Variable Request OAM PDU, is used to request one or more MIB variables from the remote peer. 	All levels
OAM Receive Statistics		
Information	Number of information PDUs received.	detail
Event	Number of loopback control PDUs received.	detail
Variable request	Number of variable request PDUs received.	detail
Variable response	Number of variable response PDUs received.	detail

Table 35: show oam ethernet link-fault-management Output Fields (*continued*)

Field Name	Field Description	Level of Output
Loopback control	Number of loopback control PDUs received.	detail
Organization specific	Number of vendor organization specific PDUs received.	detail
OAM Transmit Statistics		
Information	Number of information PDUs transmitted.	detail
Event	Number of event notification PDUs transmitted.	detail
Variable request	Number of variable request PDUs transmitted.	detail
Variable response	Number of variable response PDUs transmitted.	detail
Loopback control	Number of loopback control PDUs transmitted.	detail
Organization specific	Number of vendor organization specific PDUs transmitted.	detail
OAM Received Symbol Error Event information		
Events	Number of symbol error event TLVs that have been received after the OAM sublayer was reset.	detail
Window	Symbol error event window in the received PDU. The protocol default value is the number of symbols that can be received in one second on the underlying physical layer.	detail
Threshold	Number of errored symbols in the period required for the event to be generated.	detail
Errors in period	Number of symbol errors in the period reported in the received event PDU.	detail
Total errors	Number of errored symbols that have been reported in received event TLVs after the OAM sublayer was reset. Symbol errors are coding symbol errors.	detail
OAM Received Frame Error Event Information		
Events	Number of errored frame event TLVs that have been received after the OAM sublayer was reset.	detail
Window	Duration of the window in terms of the number of 100 ms period intervals.	detail
Threshold	Number of detected errored frames required for the event to be generated.	detail
Errors in period	Number of detected errored frames in the period.	detail

Table 35: show oam ethernet link-fault-management Output Fields (continued)

Field Name	Field Description	Level of Output
Total errors	Number of errored frames that have been reported in received event TLVs after the OAM sublayer was reset. A frame error is any frame error on the underlying physical layer.	detail
OAM Received Frame Period Error Event Information		
Events	Number of frame seconds errors event TLVs that have been received after the OAM sublayer was reset.	detail
Window	Duration of the frame seconds window.	detail
Threshold	Number of frame seconds errors in the period.	detail
Errors in period	Number of frame seconds errors in the period.	detail
Total errors	Number of frame seconds errors that have been reported in received event TLVs after the OAM sublayer was reset.	detail
OAM Transmitted Symbol Error Event Information		
Events	Number of symbol error event TLVs that have been transmitted after the OAM sublayer was reset.	detail
Window	The symbol error event window in the transmitted PDU.	detail
Threshold	Number of errored symbols in the period required for the event to be generated.	detail
Errors in period	Number of symbol errors in the period reported in the transmitted event PDU.	detail
Total errors	Number of errored symbols reported in event TLVs that have been transmitted after the OAM sublayer was reset.	detail
OAM Transmitted Frame Error Event Information		
Events	Number of errored frame event TLVs that have been transmitted after the OAM sublayer was reset.	detail
Window	Duration of the window in terms of the number of 100-ms period intervals.	detail
Threshold	Number of detected errored frames required for the event to be generated.	detail
Errors in period	Number of detected errored frames in the period.	detail
Total errors	Number of errored frames that have been detected after the OAM sublayer was reset.	detail

Sample Output

show oam ethernet link-fault-management brief

```

user@host> show oam ethernet link-fault-management brief
Interface: ge-0/0/1
Status: Running, Discovery state: Send Any
Peer address: 00:00:5E:00:53:AB
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
Remote loopback status: Disabled on local port, Enabled on peer port
Remote entity information:
  Remote MUX action: discarding, Remote parser action: loopback
  Discovery mode: active, Unidirectional mode: unsupported
  Remote loopback mode: supported, Link events: supported
  Variable requests: unsupported

```

show oam ethernet link-fault-management detail

```

user@host> show oam ethernet link-fault-management detail
Interface: ge-0/0/1
Status: Running, Discovery state: Send Any
Peer address: 00:00:5E:00:53:AC
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
OAM receive statistics:
  Information: 186365, Event: 0, Variable request: 0, Variable response: 0
  Loopback control: 0, Organization specific: 0
OAM transmit statistics:
  Information: 186347, Event: 0, Variable request: 0, Variable response: 0
  Loopback control: 0, Organization specific: 0
OAM received symbol error event information:
  Events: 0, Window: 0, Threshold: 0
  Errors in period: 0, Total errors: 0
OAM received frame error event information:
  Events: 0, Window: 0, Threshold: 0
  Errors in period: 0, Total errors: 0
OAM received frame period error event information:
  Events: 0, Window: 0, Threshold: 0
  Errors in period: 0, Total errors: 0
OAM transmitted symbol error event information:
  Events: 0, Window: 0, Threshold: 1
  Errors in period: 0, Total errors: 0
OAM transmitted frame error event information:
  Events: 0, Window: 0, Threshold: 1
  Errors in period: 0, Total errors: 0
Remote entity information:
  Remote MUX action: forwarding, Remote parser action: forwarding
  Discovery mode: active, Unidirectional mode: unsupported
  Remote loopback mode: supported, Link events: supported
  Variable requests: unsupported

```

show poe controller (View)

Supported Platforms	SRX1500, SRX320, SRX340, SRX550M
Syntax	show poe controller
Release Information	Command introduced in Junos OS Release 9.5.
Description	Display the status of the Power over Ethernet (PoE) controller.
Options	none—Display general parameters of the PoE software module controller.
Required Privilege Level	View
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring PoE on All Interfaces on page 189
Output Fields	Table 36 on page 416 lists the output fields for the show poe controller command. Output fields are listed in the approximate order in which they appear.

Table 36: show poe controller Output Fields

Field name	Field Description
Controller-index	Identifies the controller.
Maximum-power	Specifies the maximum power that can be provided by the SRX Series device to PoE ports.
Power-consumption	Specifies the total amount of power allocated to the PoE ports.
Guard-band	Shows the guard band configured on the controller.
Management	Shows the power management mode.

Sample Output

show poe controller

```
user@host>show poe controller
```

```

Controller  Maximum  Power      Guard band  Management
index      power    consumption
  0         150.0 W  0.0 W      0 W         Static

```

show pppoe interfaces

Supported Platforms	SRX1500, SRX300, SRX320, SRX340, SRX550M
Syntax	show pppoe interfaces <brief detail extensive> <pp0.logical>
Release Information	Command introduced in Junos OS Release 9.5.
Description	Display session-specific information about PPPoE interfaces.
Options	<p>none—Display interface information for all PPPoE interfaces.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>extensive—(Optional) Display information about the number of packets sent and received and the number of timeouts during a PPPoE session.</p> <p>pp0.logical—(Optional) Name of an interface. The logical unit number for static interfaces can be a value from 0 through 16,385. The logical unit number for dynamic interfaces can be a value from 1,073,741,824 through the maximum number of logical interfaces supported on your SRX300, SRX320, and SRX340, and SRX550M devices.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Understanding Ethernet Interfaces on page 117
List of Sample Output	<p>show pppoe interfaces on page 419</p> <p>show pppoe interfaces brief on page 419</p> <p>show pppoe interfaces detail on page 419</p> <p>show pppoe interfaces extensive on page 419</p>
Output Fields	Table 37 on page 417 lists the output fields for the show pppoe interfaces command. Output fields are listed in the approximate order in which they appear.

Table 37: show pppoe interfaces Output Fields

Field Name	Field Description
Index	Index number of the logical interface, which reflects its initialization sequence.
State	State of the logical interface: up or down .
Session ID	Session ID.
Service name	Type of service required (can be used to indicate an ISP name, a class, or quality of service).
Configured AC name	Configured access concentrator name.

Table 37: show pppoe interfaces Output Fields (*continued*)

Field Name	Field Description
Session AC name	Name of the access concentrator.
Remote MAC address or Remote MAC	MAC address of the remote side of the connection, either the access concentrator or the PPPoE client.
Auto-reconnect timeout	Timeout value for reconnecting after a PPPoE session is terminated (in seconds).
Idle timeout	Length of time (in seconds) that a connection can be idle before disconnecting.
Session uptime	Length of time the session has been up, in <i>hh:mm:ss</i> .
Ignore End-Of-List tag	Disables the End-of-List tag to continue processing of other tags after the End-of-List tag in a PPPoE Active Discovery Offer (PADO) packet.
Underlying interface	Interface on which PPPoE is running.
Packet Type	<p>Number of packets sent and received during the PPPoE session, categorized by packet type and packet errors:</p> <ul style="list-style-type: none"> • PADI—PPPoE Active Discovery Initiation packets. • PADO—PPPoE Active Discovery Offer packets. • PADR—PPPoE Active Discovery Request packets. • PADS—PPPoE Active Discovery Session-Confirmation packets. • PADT—PPPoE Active Discovery Termination packets. • Service name error—Packets for which the Service-Name request could not be honored. • AC system error—Packets for which the access concentrator experienced an error in performing the host request. For example, the host had insufficient resources to create a virtual circuit. • Generic error—Packets that indicate an unrecoverable error occurred. • Malformed packets—Malformed or short packets that caused the packet handler to discard the frame as unreadable. • Unknown packets—Unrecognized packets.
Timeout	<p>Timeouts that occur during the PPPoE session:</p> <ul style="list-style-type: none"> • PADI—No PADI packets received within the timeout period. • PADO—No PADO packets received within the timeout period. (This value is always zero and is not supported.) • PADR—No PADR packets received within the timeout period.
Receive Error Counters	<p>Error counters received during the PPPoE session:</p> <ul style="list-style-type: none"> • PADI—No PADI error counters received during the session. • PADO—No PADO error counters received during the session. • PADR—No PADR error counters received during the session. • PADS—No PADS error counters received during the session.

Sample Output

show pppoe interfaces

```
user@host> show pppoe interfaces
pp0.0 Index 71
  State: Session up, Session ID: 4,
  Service name: None,
  Session AC name: srx-pppoe-ac, Configured AC name: None,
  Remote MAC address: b0:c6:9a:74:5e:c1,
  Session uptime: 5d 15:21 ago,
  Auto-reconnect timeout: Never, Idle timeout: Never,
  Underlying interface: ge-0/0/1.0 Index 70
```

show pppoe interfaces brief

```
user@host> show pppoe interfaces brief
```

Interface	Underlying interface	State	Session ID	Remote MAC
pp0.0	ge-0/0/1.0	Session up	4	b0:c6:9a:74:5e:c1

show pppoe interfaces detail

```
user@host> show pppoe interfaces detail
pp0.0 Index 71
  State: Session up, Session ID: 4,
  Service name: None,
  Session AC name: srx-pppoe-ac, Configured AC name: None,
  Remote MAC address: b0:c6:9a:74:5e:c1,
  Session uptime: 5d 15:21 ago,
  Auto-reconnect timeout: Never, Idle timeout: Never,
  Underlying interface: ge-0/0/1.0 Index 70
  Ignore End-Of-List tag: Enable
```

show pppoe interfaces extensive

```
user@host> show pppoe interfaces extensive
pp0.0 Index 71
  State: Session up, Session ID: 4,
  Service name: None,
  Session AC name: srx-pppoe-ac, Configured AC name: None,
  Remote MAC address: b0:c6:9a:74:5e:c1,
  Session uptime: 5d 15:22 ago,
  Auto-reconnect timeout: Never, Idle timeout: Never,
  Underlying interface: ge-0/0/1.0 Index 70
```

PacketType	Sent	Received
PADI	1	0
PADO	0	1
PADR	1	0
PADS	0	1
PADT	0	0
Service name error	0	0
AC system error	0	0
Generic error	0	0
Malformed packets	0	0
Unknown packets	0	0

```
Timeout
  PADI      0
  PADO      0
  PADR      0
Receive Error Counters
```

PADI	0
PADO	0
PADR	0
PADS	0

show pppoe statistics

Supported Platforms [SRX1500, SRX300, SRX320, SRX340](#)

Syntax `show pppoe statistics`
`<logical-interface-name>`

Release Information Command is first introduced in Junos OS Release 9.5.

Description Display statistics information about PPPoE interfaces.

Options **none**—Display PPPoE statistics for all interfaces.
logical-interface-name—(Optional) Name of an underlying PPPoE logical interface.

Required Privilege Level view

Related Documentation

- [show pppoe interfaces on page 417](#)
- [Understanding Ethernet Interfaces on page 117](#)

List of Sample Output [show pppoe statistics on page 422](#)

Output Fields [Table 38 on page 421](#) lists the output fields for the **show pppoe statistics** command. Output fields are listed in the approximate order in which they appear.

Table 38: show pppoe statistics Output Fields

Field Name	Field Description
Active PPPoE sessions	Total number of active PPPoE sessions.
Packet Type	<p>Number of packets sent and received during the PPPoE session, categorized by packet type and packet errors:</p> <ul style="list-style-type: none"> • PADI—PPPoE Active Discovery Initiation packets. • PADO—PPPoE Active Discovery Offer packets. • PADR—PPPoE Active Discovery Request packets. • PADS—PPPoE Active Discovery Session-Confirmation packets. • PADT—PPPoE Active Discovery Termination packets. • Service name error—Packets for which the Service-Name request could not be honored. • AC system error—Packets for which the access concentrator experienced an error in performing the host request. For example, the host had insufficient resources to create a virtual circuit. • Generic error—Packets that indicate an unrecoverable error occurred. • Malformed packets—Malformed or short packets that caused the packet handler to discard the frame as unreadable. • Unknown packets—Unrecognized packets.

Table 38: show pppoe statistics Output Fields (*continued*)

Field Name	Field Description
Timeout	<p>Timeouts that occur during the PPPoE session:</p> <ul style="list-style-type: none"> • PADI—No PADI packets received within the timeout period. • PADO—No PADO packets received within the timeout period. (This value is always zero and is not supported.) • PADR—No PADR packets received within the timeout period.
Receive Error Counters	<p>Error counters received during the PPPoE session:</p> <ul style="list-style-type: none"> • PADI—No PADI error counters received during the session. • PADO—No PADO error counters received during the session. • PADR—No PADR error counters received during the session. • PADS—No PADS error counters received during the session.

Sample Output

show pppoe statistics

```

user@host> show pppoe statistics
Active PPPoE sessions: 0

PacketType          Sent      Received
PADI                0         0
PADO                0         0
PADR                0         0
PADS                0         0
PADT                0         0
Service name error  0         0
AC system error     0         0
Generic error       0         0
Malformed packets  0         0
Unknown packets     0         0
Timeout
PADI                0
PADO                0
PADR                0
Receive Error Counters
PADI                0
PADO                0
PADR                0
PADS                0
    
```

show poe telemetries

Supported Platforms	SRX1500, SRX320, SRX340, SRX550M
Syntax	show poe telemetries <interface <i>interface-name</i> count <i>number</i> > <count <i>number</i> interface <i>interface-name</i> >
Release Information	Command modified in Junos OS Release 12.3X48-D10.
Description	Display a history of power consumption on the specified interface. Telemetries must be enabled on the interface before you can display a history of power consumption.
Options	<ul style="list-style-type: none"> • Interface <i>interface-name</i>—Display telemetries for the specified PoE interface. • count <i>number</i>—Display the specified number of telemetries records for the specified PoE interface.
Required Privilege Level	View
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring PoE on All Interfaces on page 189
Output Fields	Table 39 on page 423 lists the output fields for the show poe telemetries interface command. Output fields are listed in the approximate order in which they appear.

Table 39: show poe telemetries interface Output Fields

Field name	Field Description
SI No	Number of the record for the specified port. The last record is the most recent.
Timestamp	Time that the power-consumption data was gathered.
Power	Amount of power provided by the specified port at the time the data was gathered.
Voltage	Voltage on the specified port at the time the data was gathered.

Sample Output

show poe telemetries interface

```
user@host>show poe telemetries interface ge-0/0/1 count 8
```

```

SI No   Timestamp                Power    Voltage
 1  Fri Jan 04 11:41:15 2009 6.6 W   47.2 V
 2  Fri Jan 04 11:40:15 2009 6.6 W   47.2 V
 3  Fri Jan 04 11:39:15 2009 6.6 W   47.2 V
 4  Fri Jan 04 11:38:15 2009 6.6 W   47.2 V
 5  Fri Jan 04 11:37:15 2009 6.6 W   47.2 V
 6  Fri Jan 04 11:36:15 2009 6.6 W   47.2 V

```

```
7    Fri Jan 04 11:35:15 2009 6.6 W    47.2 V
8    Fri Jan 04 11:34:15 2009 6.6 W    47.2 V
```

user@host>show poe telemetries count 5 interface ge-0/0/1

Sl No	Timestamp	Power	Voltage
1	Fri Jan 04 11:47:15 2009	6.6 W	47.2 V
2	Fri Jan 04 11:38:15 2009	6.6 W	47.2 V
3	Fri Jan 04 11:29:15 2009	6.6 W	47.2 V
4	Fri Jan 04 11:11:15 2009	6.6 W	47.2 V
5	Fri Jan 04 11:10:15 2009	6.6 W	47.2 V

show services accounting

Supported Platforms [SRX Series, vSRX](#)

Syntax show services accounting
 aggregation
 errors
 <inline-jflow | inline-jflow fpc-slot *slot number*>
 flow
 <inline-jflow | inline-jflow fpc-slot *slot number*>
 flow-detail
 memory
 packet-size-distribution
 status
 <inline-jflow | inline-jflow fpc-slot *slot number*>
 usage

Release Information Command introduced in Junos OS Release 10.4. The **inline-jflow** and **fpc-slot** options are added in Junos OS Release 12.1X45-D10.

Description Display sampled accounting service.

- Options**
- aggregation—Display aggregation information.
 - errors —Display error statistics.
 - inline-jflow — Display service accounting inline flow monitoring parameters.
 - fpc-slot *slot number*— Display Flexible PIC Concentrator (FPC) slot for inline flow monitoring.
 - flow—Display flow information.
 - inline-jflow — Display service accounting inline flow monitoring parameters.
 - fpc-slot *slot number*— Display Flexible PIC Concentrator (FPC) slot for inline flow monitoring.
 - flow-detail—Display flow detail.
 - memory—Display memory information.
 - packet-size-distribution—Display packet size distribution.
 - status—Display service accounting parameters.
 - inline-jflow — Display service accounting inline flow monitoring parameters.
 - fpc-slot *slot number*— Display Flexible PIC Concentrator (FPC) slot for inline flow monitoring.
 - usage—Display CPU usage.

Required Privilege Level view

- Related Documentation**
- [Configuring Flow Aggregation to Use Version 9 Flow Templates on page 25](#)
- List of Sample Output**
- [show services accounting status inline-jflow on page 426](#)
 - [show services accounting errors inline-jflow on page 426](#)
 - [show service accounting flow inline-jflow on page 426](#)
- Output Fields** Lists the output fields for the **show services accounting** command.

Sample Output

show services accounting status inline-jflow

```
user@host> show services accounting status inline-jflow
Status information
  FPC Slot: 5
  Export format: IP-FIX(V9)
  IPv4 Route Record Count: 16, IPv6 Route Record Count: 5
  Route Record Count: 21, AS Record Count: 1
  Route-Records Set: Yes, Config Set: Yes
```

show services accounting errors inline-jflow

```
user@host> show services accounting errors inline-jflow
Error Information
  FPC Slot: 5
  PIC Slot: 0
  Flow Creation Failures: 0
  Route Record Lookup Failures: 0
  AS Lookup Failures: 0
  Export Packet Failures: 0
  Memory Overload: No

IPv4 Errors:
  IPv4 Flow Creation Failures: 0
  IPv4 Route Record Lookup Failures: 0
  IPv4 AS Lookup Failures: 0
  IPv4 Export Packet Failures: 0

IPv6 Errors:
  IPv6 Flow Creation Failures: 0
  IPv6 Route Record Lookup Failures: 0
  IPv6 AS Lookup Failures: 0
  IPv6 Export Packet Failures: 0
```

show service accounting flow inline-jflow

```
user@host> show service accounting flow inline-jflow
Flow Information
  FPC Slot: 5
  PIC Slot: 0
  Flow Packets: 2  Flow Bytes: 0
  Active Flows: 1  Total Flows: 2
  Flows Exported: 0  Flow Packets Exported: 231
  Flows Inactive Timed Out: 1  Flows Active Timed Out: 2

IPv4 Flows:
  IPv4 Flow Packets: 1  IPv4 Flow Bytes: 0
```

IPv4 Active Flows: 1 IPv4 Total Flows: 1
IPv4 Flows Exported: 0 IPv4 Flow Packets Exported: 132
IPv4 Flows Inactive Timed Out: 0 IPv4 Flows Active Timed Out: 1

IPv6 Flows:
IPv6 Flow Packets: 1 IPv6 Flow Bytes: 0
IPv6 Active Flows: 0 IPv6 Total Flows: 1
IPv6 Flows Exported: 0 IPv6 Flow Packets Exported: 99
IPv6 Flows Inactive Timed Out: 1 IPv6 Flows Active Timed Out: 1

show services accounting aggregation (View)

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, SRX550M, vSRX](#)

Syntax show services accounting aggregation

Release Information Command introduced in Junos OS Release 10.4.

Description Display aggregation information for the accounting service.

- Options**
- as—Display aggregation type AS.
 - destination-prefix—Display aggregation type destination-prefix.
 - protocol-port—Display aggregation type protocol-port.
 - source-destination-prefix—Display aggregation type source-destination-prefix.
 - source-prefix—Display aggregation type source-prefix.
 - template—Display aggregation type template.

Required Privilege Level view

Related Documentation

- [Configuring Flow Aggregation to Use Version 9 Flow Templates on page 25](#)

show services accounting aggregation template (View)

Supported Platforms SRX1500, SRX300, SRX320, SRX340, SRX550M, vSRX

Syntax show services accounting aggregation template

Release Information Command introduced in Junos OS Release 10.4.

Description Display aggregation type template.

- Options**
- detail—Display detailed output.
 - extensive—Display extensive output.
 - template-name—Display name of the template.
 - terse—Display terse output (default).

Required Privilege Level view

Related Documentation

- [Configuring Flow Aggregation to Use Version 9 Flow Templates on page 25](#)

show services accounting flow-detail (View)

Supported Platforms SRX1500, SRX300, SRX320, SRX340, SRX550M, vSRX

Syntax show services accounting flow-detail

Release Information Command introduced in Junos OS Release 10.4.

Description Display flow detail

- Options**
- destination-as—Filter term destination AS.
 - destination-port—Filter term destination port.
 - destination-prefix—Filter term destination prefix.
 - detail—Display detailed output.
 - extensive—Display extensive output.
 - input-snmp-interface-index—Filter term input SNMP interface index.
 - limit—Display maximum number of flows to display.
 - name—Display name of the service, wildcard, or “all”.
 - order—Display order for displaying flows.
 - output-snmp-interface-index—Filter term output SNMP interface index.
 - proto—Filter term protocol.
 - source-as—Filter term source AS.

Required Privilege Level view

Related Documentation

- [Configuring Flow Aggregation to Use Version 9 Flow Templates on page 25](#)