JUNIPER NETWORKS

Junos® OS

Ethernet Interfaces Feature Guide for Security Devices

Release
12.1X46-D10

Published: 2014-07-21

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

Junos® OS Ethernet Interfaces Feature Guide for Security Devices
12.1X46-D10

Table of Contents

List of Figures

List of Tables

About the Documentation

Documentation and Release Notes

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

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

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

Supported Platforms

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

- J Series
- SRX Series
- LN Series

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

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

Documentation Conventions

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

Table 1: Notice Icons

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

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

Table 2: Text and Syntax Conventions

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

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

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

GUI Conventions

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

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

Documentation Feedback

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

- Online feedback rating system—On any page at the Juniper Networks Technical Documentation 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 https://www.juniper.net/cgi-bin/docbugreport/.
- 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 JNASC support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

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

Self-Help Online Tools and Resources

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

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

CHAPTER 1

Ethernet Interfaces

Understanding Ethernet Interfaces

Supported Platforms J Series, LN Series, SRX Series

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

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 3 on page 4 shows collision rounds up to round 10.

Table 3: Collision Backoff Algorithm Rounds

Round	Size of Set	Elements in the Set
1	2	{0,1}
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}

Table 3: Collision Backoff Algorithm Rounds *(continued)*

Round	Size of Set	Elements in the Set
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.

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 1 on page 6 shows the Ethernet frame format.

Figure 1: Ethernet Frame Format

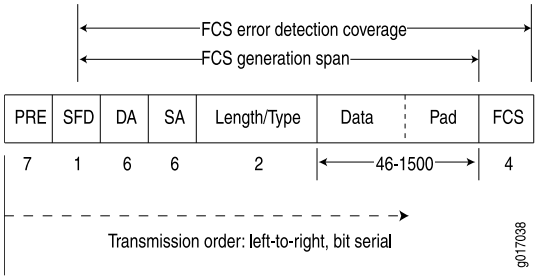

Ethernet frames have the following fields:

- 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*
- Example: Creating an Ethernet Interface on page 51
- Example: Deleting an Ethernet Interface on page 52
- Understanding Static ARP Entries on Ethernet Interfaces on page 7
- Understanding Promiscuous Mode on Ethernet Interfaces on page 7
- *Ethernet Interfaces Feature Guide for Security Devices*

Understanding Static ARP Entries on Ethernet Interfaces

Supported Platforms J Series, LN Series, SRX Series

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 3
- Example: Configuring Static ARP Entries on Ethernet Interfaces on page 53
- *Ethernet Interfaces Feature Guide for Security Devices*

Understanding Promiscuous Mode on Ethernet Interface

Supported Platforms LN Series, SRX1400, SRX3400, SRX3600, SRX5400, SRX5600, SRX5800

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.

Related Documentation

- Understanding Ethernet Interfaces on page 3
- Enabling and Disabling Promiscuous Mode on Ethernet Interfaces (CLI Procedure) on page 56
- *Ethernet Interfaces Feature Guide for Security Devices*

CHAPTER 2

Aggregated Ethernet Interfaces

Understanding Aggregated Ethernet Interface Link Speed

Supported Platforms J Series, LN Series, 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 10
- Example: Configuring Aggregated Ethernet Link Speed on page 60
- Understanding Minimum Links for Aggregated Ethernet Interfaces on page 12

- *Ethernet Interfaces Feature Guide for Security Devices*

Understanding Aggregated Ethernet Interface Removal

Supported Platforms J Series, LN Series, SRX Series

You can delete an aggregated Ethernet interface from the interface configuration. The Junos OS removes the configuration statements related to **ae*x*** 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 10
- Example: Deleting Aggregated Ethernet Interfaces on page 62
- Example: Deleting Aggregated Ethernet Interface Contents on page 63
- *Ethernet Interfaces Feature Guide for Security Devices*

Understanding Aggregated Ethernet Interfaces

Supported Platforms LN Series, SRX3400, SRX3600, SRX5400, SRX5600, SRX5800

Link aggregation of Ethernet interfaces is defined in the IEEE 802.3ad standard. The 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 and J Series devices, see the *Junos OS Layer 2 Bridging and Switching Library for Security Devices*.

This topic contains the following sections:

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, but is only supported on Layer 3.

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 3
- Aggregated Ethernet Interfaces Configuration Overview on page 57
- Understanding LACP on Standalone Devices on page 17
- Understanding LACP on Chassis Clusters on page 18
- Understanding VLAN Tagging for Aggregated Ethernet Interfaces on page 14
- Understanding Promiscuous Mode for Aggregated Ethernet Interfaces on page 13
- *Ethernet Interfaces Feature Guide for Security Devices*

Understanding Minimum Links for Aggregated Ethernet Interfaces

Supported Platforms J Series, LN Series, 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 10
- Example: Configuring Aggregated Ethernet Minimum Links on page 61
- Understanding Aggregated Ethernet Interface Link Speed on page 9
- *Ethernet Interfaces Feature Guide for Security Devices*

Understanding Physical Interfaces for Aggregated Ethernet Interfaces

Supported Platforms J Series, LN Series, 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 ae*x* 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 10
- Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces on page 59
- *Ethernet Interfaces Feature Guide for Security Devices*

Understanding Promiscuous Mode for Aggregated Ethernet Interfaces

Supported Platforms LN Series, SRX1400, SRX3400, SRX3600, 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 10
- Aggregated Ethernet Interfaces Configuration Overview on page 57
- *Ethernet Interfaces Feature Guide for Security Devices*

Understanding the Aggregated Ethernet Interfaces Device Count

Supported Platforms J Series, LN Series, 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 10
- Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device on page 58
- Example: Deleting Aggregated Ethernet Interfaces on page 62
- *Ethernet Interfaces Feature Guide for Security Devices*

Understanding VLAN Tagging for Aggregated Ethernet Interfaces

Supported Platforms LN Series, SRX1400, SRX3400, SRX3600, 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 10
- Aggregated Ethernet Interfaces Configuration Overview on page 57
- *Ethernet Interfaces Feature Guide for Security Devices*

CHAPTER 3

Link Aggregation Control Protocol

Understanding LACP on Standalone Devices

Supported Platforms J Series, LN Series, 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).

Related Documentation

- Understanding Aggregated Ethernet Interfaces on page 10
- Understanding LACP on Chassis Clusters on page 18
- Example: Configuring LACP on Standalone Devices on page 65
- *Ethernet Interfaces Feature Guide for Security Devices*

Understanding LACP on Chassis Clusters

Supported Platforms J Series, LN Series, SRX Series

Link aggregation groups (LAGs) can be established across nodes in a chassis cluster.

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 on the other node. You can configure up to eight active links and eight standby links per node. Link aggregation allows a redundant Ethernet interface to add multiple child interfaces from both nodes and thereby create 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.

Standalone LAG interfaces are supported on clustered devices but cannot be added to redundant Ethernet interfaces. Instead, aggregated Ethernet interfaces and redundant Ethernet interfaces coexist. (Because the functionality of a redundant Ethernet interface relies on the Junos OS aggregated Ethernet framework, you can think of it as a special aggregated Ethernet interface.) Likewise, any child interface of an existing LAG cannot be added to a redundant Ethernet interface and vice versa. The maximum number of total combined standalone aggregate interfaces (**ae**) and redundant Ethernet interfaces (**reth**) per cluster is 128.

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:

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 will be down even if some links are still working.

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 will be 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.

Hitless Failover

With LACP, it is essential for the redundant Ethernet interface to support hitless failover between the active and standby links in normal operation. The term *hitless* means that the redundant Ethernet interface state remains up during failover.

The lacpd process manages both the active and standby links of the redundant Ethernet interfaces. A redundant Ethernet interface pseudolink is in the up condition when the number of active up links is not less than the number of minimum links configured. Therefore, to support hitless failover, the LACP state on the redundant Ethernet interface standby links must be collecting and distributing before failover occurs.

PDUs

By default, aggregated and redundant Ethernet links do not exchange link aggregation control protocol data units (PDUs), which contain information about the state of the link. You can configure Ethernet links to actively transmit link aggregation control PDUs, or you can configure the links to passively transmit them, sending out link aggregation control PDUs only when they receive them 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).

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

- Understanding Aggregated Ethernet Interfaces on page 10
- Understanding LACP on Standalone Devices on page 17
- Example: Configuring LACP on Chassis Clusters on page 66
- *Ethernet Interfaces Feature Guide for Security Devices*

LAG and LACP Support on the SRX5000 Module Port Concentrator

Supported Platforms SRX5600, SRX5800

The SRX5000 Module Port Concentrator (SRX5K-MPC) on 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 balance (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.

Related Documentation

- Aggregated Ethernet Interfaces Configuration Overview on page 57
- Example: Configuring LACP on Standalone Devices on page 65
- Example: Configuring LACP on Chassis Clusters on page 66
- *Ethernet Interfaces Feature Guide for Security Devices*

CHAPTER 4

Gigabit Ethernet Physical Interface Modules

Understanding the 1-Port Gigabit Ethernet SFP Mini-PIM

Supported Platforms LN Series, SRX100, SRX210, SRX220, SRX240

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

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: 1-Port Gigabit Ethernet SFP mini-PIM does not support family ethernet-switching.

Related Documentation

- Understanding Ethernet Interfaces on page 3
- Example: Configuring the 1-Port Gigabit Ethernet SFP Mini-PIM Interface on page 69
- *Ethernet Interfaces Feature Guide for Security Devices*

Understanding the 8-Port Gigabit Ethernet SFP XPIM

Supported Platforms LN Series, SRX550, SRX650

A Gigabit Ethernet Physical Interface Module (XPIM) is a network interface card (NIC) that installs in the front slots of the SRX550 or SRX650 Services Gateway to provide physical connections to a LAN or a WAN.

Small form-factor pluggables (SFPs) are hot-pluggable modular interface transceivers for gigabit and Fast Ethernet connections. The 8-port SFP Gigabit Ethernet interface enables customers to connect to Ethernet WAN services as well as to local servers at gigabit speed.

Supported Features

The following features are supported on the 8-Port Gigabit Ethernet SFP XPIM:

- Operates on both a slot with a maximum bandwidth of 8 gigabits and a slot with a maximum bandwidth of 1 gigabit
- Operates in tri-rate (10/100/1000 Mbps) mode with copper SFPs

- Routing and switched mode operation
- Layer 2 protocols
  - Link Aggregation Control Protocol (LACP)
  - Link Layer Discovery Protocol (LLDP)
  - GARP VLAN Registration Protocol (GVRP)
  - Internet Group Management Protocol (IGMP) snooping (v1 and v2)
  - Spanning Tree Protocol (STP), Real-Time Streaming Protocol (RTSP), and Multiple Spanning Tree Protocol (MSTP)
  - 802.1x
- Encapsulation (supported at the Physical Layer)
  - ethernet-bridge
  - ethernet-ccc
  - ethernet-tcc
  - ethernet-vpls
  - extended-vlan-ccc
  - extended-vlan-tcc
  - flexible-ethernet-services
  - vlan-ccc
- Q in Q VLAN tagging
- Integrated routing and bridging (IRB)
- Jumbo frames (9192 byte size)
- Chassis cluster switching
- Chassis cluster fabric link using GE ports

NOTE:

The following Layer 2 switching features are not supported when the 8-Port Gigabit Ethernet SFP XPIM is plugged in slots with speeds of less than 1 gigabit:

- Q in Q VLAN tagging
- Link aggregation using ports across multiple XPIMs

Interface Names and Settings

The following format is used to represent the 8-Port SFP XPIM:

type-fpc/pic/port

Where:

- type—Media type (ge)
- fpc—Number of the Flexible PIC Concentrator (FPC) card where the physical interface resides
- pic—Number of the PIC where the physical interface resides (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 maximum transmission unit (MTU) size for the XPIM. Junos OS supports values from 256 through 9192. The default MTU size for the 8-Port Gigabit Ethernet SFP XPIM is 1514.

Related Documentation

- Example: Configuring 8-Port Gigabit Ethernet SFP XPIMs on page 74
- *Ethernet Interfaces Feature Guide for Security Devices*

Understanding the 2-Port 10-Gigabit Ethernet XPIM

Supported Platforms LN Series, SRX650

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 (1Gigabit 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 auto negotiation is not supported.**

This topic includes the following sections:

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 3
- Example: Configuring the 2-Port 10-Gigabit Ethernet XPIM Interface on page 89
- *Ethernet Interfaces Feature Guide for Security Devices*

CHAPTER 5

Ethernet OAM Link Fault Management

Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways

Supported Platforms LN Series, SRX100, SRX210, SRX220, SRX240, SRX550, SRX650

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.

This feature is supported on SRX100, SRX210, SRX220, SRX240, SRX550, and SRX650 devices.

NOTE: For SRX550 and SRX650 devices, LFM is supported only on devices that have 16-port or 24-port GPIMs.

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 4 on page 32 lists the interfaces modes supported.

Table 4: Supported Interface Modes

Interfaces	Mode
Physical interface (fe/ge)	Family • ccc • ethernet-switching • inet6 • inet • iso • mpls • tcc
	IFD encapsulations • ethernet-ccc • extended-vlan-ccc (IFD vlan-tagging mode) • ethernet-tcc • extended-vlan-tcc

Table 4: Supported Interface Modes *(continued)*

Interfaces	Mode
Aggregated Ethernet interface (Static or LACP lag)	Family • ethernet-switching • inet • mpls • iso • inet6
	IFD encapsulations • ethernet-ccc • extended-vlan-ccc (IFD vlan-tagging mode) • vlan-ccc

Related Documentation

- *Ethernet Port Switching Feature Guide for Security Devices*
- Example: Configuring Ethernet OAM Link Fault Management on page 95
- *Ethernet Interfaces Feature Guide for Security Devices*

CHAPTER 6

Power over Ethernet

Understanding Power over Ethernet

Supported Platforms LN Series, SRX210, SRX240, SRX650

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

Table 5 on page 35 lists the PoE specifications for the SRX210, SRX240, and SRX650 devices

Table 5: PoE Specifications for the SRX210, SRX240 and SRX650 Devices

Specifications	For SRX210 Device	For SRX240 Device	For SRX650 Device
Supported standards	• IEEE 802.3 AF • Legacy (pre-standards)	• IEEE 802.3 AF • IEEE 802.3 AT (PoE+) • Legacy (pre-standards)	• IEEE 802.3 AF • IEEE 802.3 AT (PoE+) • Legacy (pre-standards)

Table 5: PoE Specifications for the SRX210, SRX240 and SRX650 Devices *(continued)*

Specifications	For SRX210 Device	For SRX240 Device	For SRX650 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**).	Supported on the following ports:
			• Slot 2 or 6 on 16 Gigabit Ethernet ports
			  ◦ **ge-2/0/0 to ge-2/0/15**
			  ◦ **ge-6/0/0 to ge-6/0/15**
			• Slot 2 or 6 on 24 Gigabit Ethernet ports
			  ◦ **ge-2/0/0 to ge-2/0/23**
			  ◦ **ge-6/0/0 to ge-6/0/23**

Total PoE power sourcing capacity	50 W	150 W	The 645 watts AC and 645 watts DC power supplies support the following capacities: • 250 watts on a single power supply, or with redundancy using the two-power-supply option. • 500 watts with the two-power-supply option operating as nonredundant.
Default per port power limit	15.4 W	15.4 W	15.4 W
Maximum per port power limit	30 W	30 W	30 W
Power management modes	• Static: Power allocated for each interface can be configured. • Class: Power allocated for interfaces is based on the class of powered device connected.	• Static: Power allocated for each interface can be configured. • Class: Power allocated for interfaces is based on the class of powered device connected.	• Static: Power allocated for each interface can be configured. • Class: Power allocated for interfaces is based on the class of powered device connected.

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 6 on page 37 lists the classes and their power ratings as specified by the IEEE standards.

Table 6: 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 3
- Example: Configuring PoE on All Interfaces on page 101
- Example: Configuring PoE on an Individual Interface on page 103
- Example: Disabling a PoE Interface on page 106
- *Ethernet Interfaces Feature Guide for Security Devices*

CHAPTER 7

Interface Properties

MTU Default and Maximum Values

Supported Platforms J Series, LN Series, 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. Refer Table 7 on page 39 and Table 8 on page 41 for default MTU values.

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

Table 7 on page 39 lists MTU values for J Series devices.

Table 7: MTU Values for J4350 and J6350 Interfaces

J4350 and J6350 Interfaces	Default Media MTU (bytes)	Maximum MTU (bytes)	Default IP MTU (bytes)
Gigabit Ethernet (10/100/1000) built-in interface	1514	9018	1500
6-Port, 8-Port, and 16-Port Gigabit Ethernet uPIMs	1514	9014	1500
Gigabit Ethernet (10/100/1000) ePIM	1514	9018	1500
Gigabit Ethernet (10/100/1000) SFP ePIM	1514	9018	1500
4-Port Fast Ethernet (10/100) ePIM	1514	1514	1500
Dual-Port Fast Ethernet (10/100) PIM	1514	9192	1500

Table 7: MTU Values for J4350 and J6350 Interfaces *(continued)*

J4350 and J6350 Interfaces	Default Media MTU (bytes)	Maximum MTU (bytes)	Default IP MTU (bytes)
Dual-Port Serial PIM	1504	9150	1500
Dual-Port T1 or E1 PIM	1504	9192	1500
Dual-Port Channelized T1/E1/ISDN PRI PIM (channelized to DS0s)	1504	4500	1500
Dual-Port Channelized T1/E1/ISDN PRI PIM (clear-channel T1 or E1)	1504	9150	1500
Dual-Port Channelized T1/E1/ISDN PRI PIM (ISDN PRI dialer interface)	1504	4098	1500
T3 (DS3) or E3 PIM	4474	9192	4470
4-Port ISDN BRI PIM	1504	4092	1500
ADSL2+ PIM (Encapsulation)			
atm-snap	4482	9150	4470
atm-vcmux	4482	9150	4470
atm-nlpid	4482	9150	4470
atm-cisco-nlpid	4482	9150	4470
ether-over-atm-llc	4482	9150	1500
atm-ppp-llc	4482	9150	4470
atm-ppp-vcmux	4482	9150	4470
atm-mlppp-llc	4482	9150	1500
ppp-over-ether-over-atm-llc	4482	9150	1492
G.SHDSL PIM (Encapsulation)			
atm-snap	4482	9150	4470
atm-vcmux	4482	9150	4470
atm-nlpid	4482	9150	4470

Table 7: MTU Values for J4350 and J6350 Interfaces *(continued)*

J4350 and J6350 Interfaces	Default Media MTU (bytes)	Maximum MTU (bytes)	Default IP MTU (bytes)
atm-cisco-nlpid	4482	9150	4470
ether-over-atm-llc	4482	9150	1500
atm-ppp-llc	4482	9150	4470
atm-ppp-vcmux	4482	9150	4470
atm-mlppp-llc	4482	9150	1500
ppp-over-ether-over-atm-llc	4482	9150	1492

Table 8 on page 41 lists MTU values for the SRX Series Services Gateways Physical Interface Modules (PIMs).

Table 8: 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)			

Table 8: MTU Values for the SRX Series Services Gateways PIMs *(continued)*

PIM	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP MTU (Bytes)
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
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

Table 8: MTU Values for the SRX Series Services Gateways PIMs *(continued)*

PIM	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP MTU (Bytes)
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

- *Junos OS Interfaces Library for Security Devices*
- Understanding Interface Physical Properties on page 43

Understanding Interface Physical Properties

Supported Platforms J Series, LN Series, 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 43 summarizes some key physical properties of device interfaces.

Table 9: Interface Physical Properties

Physical Property	Description
bert-error-rate	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*.

Table 9: Interface Physical Properties *(continued)*

Physical Property	Description
bert-period	Bit error rate test (BERT) time period over which bit errors are sampled. See *Understanding Bit Error Rate Testing*.
chap	Challenge Handshake Authentication Protocol (CHAP). Specifying **chap** enables CHAP authentication on the interface. See *Understanding CHAP Authentication on a PPPoE Interface*.
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*.
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*.
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 39.
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*.
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

- *Interfaces Overview Feature Guide for Security Devices*
- *Understanding Interfaces*
- *Understanding Bit Error Rate Testing*
- *Understanding Interface Clocking*
- *Understanding Frame Check Sequences*
- MTU Default and Maximum Values on page 39

Understanding Interface Logical Properties

Supported Platforms J Series, LN Series, SRX Series

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

- *Interfaces Overview Feature Guide for Security Devices*
- *Understanding Interfaces*
- *Understanding Protocol Families*
- *Understanding IPv4 Addressing*
- *Understanding IPv6 Addressing*
- *Understanding Virtual LANs*

Interface Naming Conventions

Supported Platforms J Series, LN Series, SRX Series

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 with 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 10 on page 46.

Table 10: Network Interface Names

Name Part	Meaning	Possible Values
type	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.
slot	Number of the chassis slot in which a PIM or IOC is installed.	J Series Services Router: The slot number begins at **1** and increases as follows from top to bottom, left to right: • J2320 router—Slots 1 to 3 • J2350 router—Slots 1 to 5 • J4350 or J6350 router—PIM slots 1 to 6 The slot number **0** is reserved for the out-of-band management ports. SRX5600 and SRX5800 devices: The slot number begins at **0** and increases as follows from left to right, bottom to top: • SRX5600 device—Slots 0 to 5 • SRX5800 device—Slots 0 to 5, 7 to 11 SRX3400 and SRX3600 devices: The Switch Fabric Board (SFB) is always **0**. Slot numbers increase as follows from top to bottom, left to right: • SRX3400 devce—Slots 0 to 4 • SRX3600 device—Slots 0 to 6
pim-or-ioc	Number of the PIM or IOC on which the physical interface is located.	J Series devices: This number is always **0**. Only one PIM can be installed in a slot. 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**. SRX3400 and SRX3600 devices: This number is always **0**. Only one IOC can be installed in a slot.

Table 10: Network Interface Names *(continued)*

Name Part	Meaning	Possible Values
port	Number of the port on a PIM or IOC on which the physical interface is located.	J Series Services Routers: • On a single-port PIM, the number is always **0**. • On a multiple-port PIM, this number begins at **0** and increases from left to right, bottom to top, to a maximum of **3**. On SRX5600 and SRX5800 devices: • 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**. On SRX3400 and SRX3600 devices: • For the SFB built-in copper Gigabit Ethernet ports, this number begins at **0** and increases from top to bottom, left to right, to a maximum of **7**. For the SFB built-in fiber Gigabit Ethernet ports, this number begins at **8** and increases from left to right to a maximum of **11**. • For 16-port Gigabit Ethernet IOCs, this number begins at **0** to a maximum of **15**. • For 2-port 10-Gigabit Ethernet IOCs, this number is **0** or **1**. Port numbers appear on the PIM or IOC faceplate.
channel	Number of the channel (time slot) on a fractional or channelized T1 or E1 interface.	• 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**.
unit	Number of the logical interface created on a physical interface.	A value from **0** through **16384**. If no logical interface number is specified, unit **0** is the default, but must be explicitly configured. NOTE: A VoIP interface must have the logical interface number 0.

For example, the interface name **e1-5/0/0:15.0** on a J Series Services Router represents the following information:

- E1 WAN interface
- PIM slot **5**
- PIM number **0** (always **0**)
- Port **0**
- Channel **15**
- Logical interface, or unit, **0**

Related Documentation

- *Understanding Interfaces*
- *Interfaces Overview Feature Guide for Security Devices*

PART 2

Configuration

CHAPTER 8

Ethernet Interfaces

Example: Creating an Ethernet Interface

Supported Platforms J Series, LN Series, SRX Series

This example shows how to create an Ethernet interface.

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 3
- Example: Deleting an Ethernet Interface on page 52
- *Ethernet Interfaces Feature Guide for Security Devices*

Example: Deleting an Ethernet Interface

Supported Platforms J Series, LN Series, SRX Series

This example shows how to delete an Ethernet interface.

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 3
- Example: Creating an Ethernet Interface on page 51
- *Ethernet Interfaces Feature Guide for Security Devices*

Example: Configuring Static ARP Entries on Ethernet Interfaces

Supported Platforms J Series, LN Series, SRX Series

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 commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set interfaces 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

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 7
- *Ethernet Interfaces Feature Guide for Security Devices*

Enabling and Disabling Promiscuous Mode on Ethernet Interfaces (CLI Procedure)

Supported Platforms LN Series, SRX1400, SRX3400, SRX3600, SRX5400, 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 Interfaces on page 7
- Understanding Ethernet Interfaces on page 3
- *Ethernet Interfaces Feature Guide for Security Devices*

CHAPTER 9

Aggregated Ethernet Interfaces

Aggregated Ethernet Interfaces Configuration Overview

Supported Platforms LN Series, SRX3400, SRX3600, SRX5400, SRX5600, SRX5800

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 58.
2. Associate a physical interface with the aggregated Ethernet interface. See "Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces" on page 59.
3. (Optional) Set the required link speed for all the interfaces included in the bundle. See "Example: Configuring Aggregated Ethernet Link Speed" on page 60.
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 61.
5. (Optional) Enable or disable VLAN tagging. See "Understanding VLAN Tagging for Aggregated Ethernet Interfaces" on page 14.
6. (Optional) Enable promiscuous mode. See "Understanding Promiscuous Mode for Aggregated Ethernet Interfaces" on page 13.

Related Documentation

- *Junos OS Layer 2 Bridging and Switching Library for Security Devices*
- Understanding Aggregated Ethernet Interfaces on page 10
- Example: Configuring LACP on Standalone Devices on page 65
- Example: Configuring LACP on Chassis Clusters on page 66
- *Ethernet Interfaces Feature Guide for Security Devices*

Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device

Supported Platforms J Series, LN Series, SRX Series

This example shows how to configure the number of aggregated Ethernet interfaces on a device.

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 13
- Aggregated Ethernet Interfaces Configuration Overview on page 57

- Example: Deleting Aggregated Ethernet Interfaces on page 62
- Verifying Aggregated Ethernet Interfaces on page 157
- *Ethernet Interfaces Feature Guide for Security Devices*

Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces

Supported Platforms J Series, LN Series, SRX Series

This example shows how to associate physical interfaces with aggregated Ethernet interfaces.

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 58.

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 13
- Aggregated Ethernet Interfaces Configuration Overview on page 57
- Verifying Aggregated Ethernet Interfaces on page 157
- *Ethernet Interfaces Feature Guide for Security Devices*

Example: Configuring Aggregated Ethernet Link Speed

Supported Platforms J Series, LN Series, SRX Series

This example shows how to configure the aggregated Ethernet link speed.

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 58.
- Associate physical interfaces with the aggregated Ethernet Interfaces. See "Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces" on page 59.

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 9
- Aggregated Ethernet Interfaces Configuration Overview on page 57
- Verifying Aggregated Ethernet Interfaces on page 157
- *Ethernet Interfaces Feature Guide for Security Devices*

Example: Configuring Aggregated Ethernet Minimum Links

Supported Platforms J Series, LN Series, 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

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 58.
- Associate physical interfaces with the aggregated Ethernet Interfaces. See "Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces" on page 59.
- Configure the aggregated Ethernet link speed. See "Example: Configuring Aggregated Ethernet Link Speed" on page 60.

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 9
- Aggregated Ethernet Interfaces Configuration Overview on page 57
- Verifying Aggregated Ethernet Interfaces on page 157
- *Ethernet Interfaces Feature Guide for Security Devices*

Example: Deleting Aggregated Ethernet Interfaces

Supported Platforms J Series, LN Series, SRX Series

This example shows how to delete aggregated Ethernet interfaces using the device count.

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 58.

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 57
- Example: Deleting Aggregated Ethernet Interface Contents on page 63
- Verifying Aggregated Ethernet Interfaces on page 157
- *Ethernet Interfaces Feature Guide for Security Devices*

Example: Deleting Aggregated Ethernet Interface Contents

Supported Platforms J Series, LN Series, SRX Series

This example shows how to delete the contents of an aggregated Ethernet interface.

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 58.
- Associate a physical interface with the aggregated Ethernet interface. See "Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces" on page 59.
- Set the required link speed for all the interfaces included in the bundle. See "Example: Configuring Aggregated Ethernet Link Speed" on page 60.
- 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 61.

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 57
- Example: Deleting Aggregated Ethernet Interfaces on page 62
- Verifying Aggregated Ethernet Interfaces on page 157
- *Ethernet Interfaces Feature Guide for Security Devices*

CHAPTER 10

Link Aggregation Control Protocol

Example: Configuring LACP on Standalone Devices

Supported Platforms J Series, LN Series, SRX Series

This example shows how to configure LACP on standalone devices.

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 58.
- Associate physical interfaces with the aggregated Ethernet Interfaces. See "Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces" on page 59.
- Configure the aggregated Ethernet link speed. See "Example: Configuring Aggregated Ethernet Link Speed" on page 60.
- Configure the aggregated Ethernet minimum links speed. See "Example: Configuring Aggregated Ethernet Minimum Links" on page 61.

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

- *Ethernet Interfaces Feature Guide for Security Devices*

Example: Configuring LACP on Chassis Clusters

Supported Platforms J Series, LN Series, SRX Series

This example shows how to configure LACP on chassis clusters.

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 58.
- Associate physical interfaces with the aggregated Ethernet Interfaces. See "Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces" on page 59.
- Configure the aggregated Ethernet link speed. See "Example: Configuring Aggregated Ethernet Link Speed" on page 60.

- Configure the aggregated Ethernet minimum links speed. See "Example: Configuring Aggregated Ethernet Minimum Links" on page 61.
- Configure the LACP on standalone devices. See "Example: Configuring LACP on Standalone Devices" on page 65.

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 *Using the CLI Editor in Configuration Mode*.

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

Confirm that the configuration is working properly.

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 18
- Verifying LACP on Standalone Devices on page 160
- *Ethernet Interfaces Feature Guide for Security Devices*

CHAPTER 11

Gigabit Ethernet Physical Interface Modules

Example: Configuring the 1-Port Gigabit Ethernet SFP Mini-PIM Interface

Supported Platforms LN Series, SRX100, SRX210, SRX220, SRX240

This example shows how to perform basic configuration for the 1-Port Gigabit Ethernet SFP Mini-PIM.

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 51.

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 commands, paste them into a text file, remove any line breaks, change any details necessary to match your network

configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set interfaces ge-2/0/0 link-mode full-duplex speed 100m
set interface ge-2/0/0 gigether-options no-loopback
```

GUI Step-by-Step Procedure

Configuring Physical Properties—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—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—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-2/0/0.0,** 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—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—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**.

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 gigether-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;
  gigether-options {
  no-loopback;
}
unit 0 {
  family inet {
  14.1.1.1/24
  }
}
```

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 72
- Verifying the FPC Status on page 73
- Verifying the Interface Settings on page 73

Verifying That the Correct Hardware Is Installed

Purpose Verify that the 1-Port Gigabit Ethernet SFP Mini-PIM is installed on the device.

Action From operational mode, enter the **show chassis hardware** command.

```
user@host> show chassis hardware detail
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                AG0309AA0004      SRX240b
Routing Engine   REV 16   750-021792   VL3180            RE-SRX240B
  da0     999 MB  ST72682                                Nand Flash
  usb0 (addr 1)  DWC OTG root hub 0    vendor 0x0000     uhub0
  usb0 (addr 2)  product 0x005a 90  vendor 0x0409     uhub1
  usb0 (addr 3)  ST72682  High Speed Mode 64218 STMicroelectronics umass0
FPC 0                                                  FPC
  PIC 0                                                16x GE Base PIC
FPC 1                    750-023367   112009000278      FPC
  PIC 0                                                1x T1E1 mPIM
FPC 2            REV 00   750-03273   AABC5081          FPC
  PIC 0                                                1x GE High-Perf SFP mPIM

    Xcvr 0       REV 02   740-011612   9101465           SFP-T
FPC 4                    750-029145   122009000061      FPC
PIC 0                                                  1x GE SFP mPIM
    Xcvr 0       REV 01   740-011782   PBL0C3T           SFP-SX
Power Supply 0
```

Verify that the output contains the following values:

- **FPC 2, PIC 0** —1x GE High-Perf SFP mPIM
- **FPC 4, PIC 0** —1x GE SFP mPIM

NOTE: In the example shown above, the output for 1-Port SFP Mini-Physical Interface Module is displayed as 1X GE SFP mPIM and the output for 1-Port Gigabit Ethernet SFP Mini-Physical Interface Module is displayed as 1X GE High-Perf SFP mPIM.

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
                     Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State           (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: Sendbcast-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 3
- Understanding the 1-Port Gigabit Ethernet SFP Mini-PIM on page 23
- *Example: Configuring the Device as a DHCP Client*
- *Ethernet Interfaces Feature Guide for Security Devices*

Example: Configuring 8-Port Gigabit Ethernet SFP XPIMs

Supported Platforms LN Series, SRX550, SRX650

This example shows how to perform a basic back-to-back device configuration with 8-port Gigabit Ethernet small form-factor pluggable (SFP) XPIMs. It describes a common scenario in which SFP XPIMs are deployed.

Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.1X44-D10 or later for SRX Series Services Gateways.
- Two SRX650 devices connected back-to-back.
- Two 8-port Gigabit Ethernet SFP XPIMs.
- Eight pairs of SFP transceivers as mentioned in *8-Port Gigabit Ethernet SFP XPIM Supported Modules* and eight cables to connect them.

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 51.

Overview and Topology

In this example, you configure two SRX650 devices. On each device you configure eight interfaces (ge-6/0/0 through ge-6/0/7), set the maximum transmission unit (MTU) value to 9192, and define a logical interface that you can connect to the 8-port SFP XPIM.

Figure 2 on page 76 shows the topology used in this example.

Figure 2: Basic Back-to-Back Device Configuration

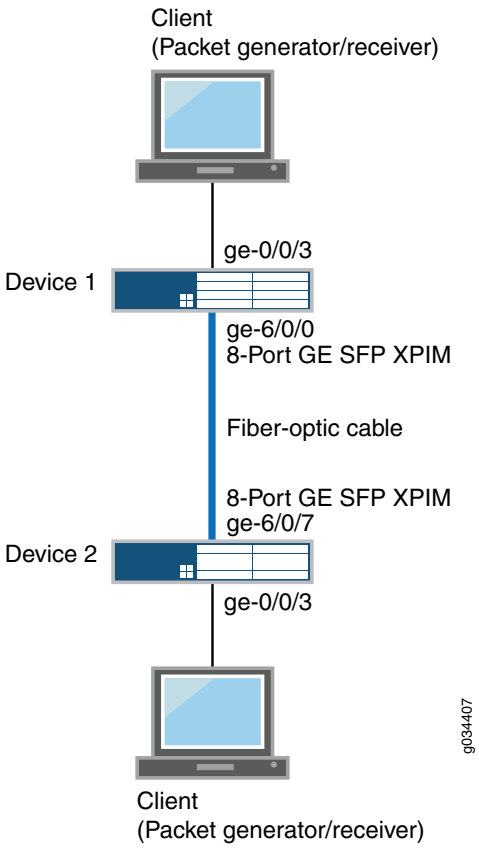

Configuration

CLI Quick Configuration

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

Device 1

```
set interfaces ge-6/0/0 mtu 9192
set interfaces ge-6/0/0 unit 0 family inet address 10.1.1.1/24
set interfaces ge-6/0/1 mtu 9192
set interfaces ge-6/0/1 unit 0 family inet address 11.1.1.1/24
set interfaces ge-6/0/2 mtu 9192
set interfaces ge-6/0/2 unit 0 family inet address 12.1.1.1/24
set interfaces ge-6/0/3 mtu 9192
set interfaces ge-6/0/3 unit 0 family inet address 13.1.1.1/24
set interfaces ge-6/0/4 mtu 9192
set interfaces ge-6/0/4 unit 0 family inet address 14.1.1.1/24
set interfaces ge-6/0/5 mtu 9192
set interfaces ge-6/0/5 unit 0 family inet address 15.1.1.1/24
set interfaces ge-6/0/6 mtu 9192
set interfaces ge-6/0/6 unit 0 family inet address 16.1.1.1/24
set interfaces ge-6/0/7 mtu 9192
set interfaces ge-6/0/7 unit 0 family inet address 17.1.1.1/24
```

Device 2

```
set interfaces ge-6/0/0 mtu 9192
set interfaces ge-6/0/0 unit 0 family inet address 10.1.1.2/24
set interfaces ge-6/0/1 mtu 9192
set interfaces ge-6/0/1 unit 0 family inet address 11.1.1.2/24
set interfaces ge-6/0/2 mtu 9192
set interfaces ge-6/0/2 unit 0 family inet address 12.1.1.2/24
set interfaces ge-6/0/3 mtu 9192
set interfaces ge-6/0/3 unit 0 family inet address 13.1.1.2/24
set interfaces ge-6/0/4 mtu 9192
set interfaces ge-6/0/4 unit 0 family inet address 14.1.1.2/24
set interfaces ge-6/0/5 mtu 9192
set interfaces ge-6/0/5 unit 0 family inet address 15.1.1.2/24
set interfaces ge-6/0/6 mtu 9192
set interfaces ge-6/0/6 unit 0 family inet address 16.1.1.2/24
set interfaces ge-6/0/7 mtu 9192
set interfaces ge-6/0/7 unit 0 family inet address 17.1.1.2/24
```

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 the interfaces on Device 1:

1. Configure the interface.

   ```
   [edit]
   user@host# set interfaces ge-6/0/0
   ```

2. Assign the maximum transmission unit value for the interface.

   ```
   [edit interfaces ge-6/0/0]
   user@host# set mtu 9192
   ```

3. Add the logical interface.

   ```
   [edit interfaces ge-6/0/0]
   user@host# set unit 0 family inet address 10.1.1.1/24
   ```

NOTE: Repeat these steps for the remaining seven ports on Device 1.

Step-by-Step Procedure

To configure the interfaces on Device 2:

1. Configure the interface.

   ```
   [edit]
   user@host# edit interfaces ge-6/0/0
   ```

2. Assign the maximum transmission unit value for the interface.

   ```
   [edit interfaces ge-6/0/0]
   user@host# set mtu 9192
   ```

3. Add the logical interface.

   ```
   [edit interfaces ge-6/0/0]
   user@host# set unit 0 family inet address 10.1.1.2/24
   ```

NOTE: Repeat these steps for the remaining seven ports on Device 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.

Device 1

```
[edit]
user@host# show interfaces
ge-6/0/0 {
   mtu 9192;
   unit 0 {
      family inet {
         address 10.1.1.1/24;
      }
   }
}
ge-6/0/1 {
   mtu 9192;
   unit 0 {
      family inet {
         address 11.1.1.1/24;
      }
   }
}
ge-6/0/2 {
   mtu 9192;
   unit 0 {
      family inet {
         address 12.1.1.1/24;
      }
   }
}
ge-6/0/3 {
   mtu 9192;
   unit 0 {
      family inet {
         address 13.1.1.1/24;
      }
   }
}
ge-6/0/4 {
   mtu 9192;
   unit 0 {
      family inet {
         address 14.1.1.1/24;
      }
   }
}
ge-6/0/5 {
   mtu 9192;
   unit 0 {
      family inet {
```

```
        address 15.1.1.1/24;
      }
    }
  }
  ge-6/0/6 {
    mtu 9192;
    unit 0 {
      family inet {
        address 16.1.1.1/24;
      }
    }
  }
  ge-6/0/7 {
    mtu 9192;
    unit 0 {
      family inet {
        address 17.1.1.1/24;
      }
    }
  }
```

Device 2

```
[edit]
user@host# show interfaces
ge-6/0/0 {
  mtu 9192;
  unit 0 {
    family inet {
      address 10.1.1.2/24;
    }
  }
}
ge-6/0/1 {
  mtu 9192;
  unit 0 {
    family inet {
      address 11.1.1.2/24;
    }
  }
}
ge-6/0/2 {
  mtu 9192;
  unit 0 {
    family inet {
      address 12.1.1.2/24;
    }
  }
}
ge-6/0/3 {
  mtu 9192;
  unit 0 {
    family inet {
      address 13.1.1.2/24;
    }
  }
}
ge-6/0/4 {
```

```
    mtu 9192;
    unit 0 {
      family inet {
        address 14.1.1.2/24;
      }
    }
  }
  ge-6/0/5 {
    mtu 9192;
    unit 0 {
      family inet {
        address 15.1.1.2/24;
      }
    }
  }
  ge-6/0/6 {
    mtu 9192;
    unit 0 {
      family inet {
        address 16.1.1.2/24;
      }
    }
  }
  ge-6/0/7 {
    mtu 9192;
    unit 0 {
      family inet {
        address 17.1.1.2/24;
      }
    }
  }
```

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

Verification

Confirm that the configuration is working properly.

Verifying the Hardware was Properly Installed

Purpose Verify that the 8-Port Gigabit Ethernet SFP XPIM is installed on the device.

Action From operational mode, enter the **show chassis hardware** command.

```
user@host> show chassis hardware detail
Hardware inventory:
```

```
Item             Version  Part number  Serial number     Description
Chassis                                AJ3009AA0001      SRX650
Midplane         REV 08   710-023875   AAAK0059
System IO        REV 08   710-023209   AAAJ9290          SRXSME System IO
Routing Engine   REV 13   750-023223   AAAJ1987          RE-SRXSME-SRE6
  ad0    2000 MB  CF 2GB               2009A     0000194075 Compact Flash
  usb0 (addr 1)  DWC OTG root hub 0    vendor 0x0000     uhub0
  usb0 (addr 2)  product 0x005a 90     vendor 0x0409     uhub1
FPC 0                                                   FPC
  PIC 0                                                 4x GE Base PIC
FPC 1            REV 03   750-038290   AADL2016          FPC
FPC 5                                                   FPC
  PIC 0                                                 8x GE SFP gPIM
FPC 6            REV 03   750-037551   AAEC8065          FPC
  PIC 0                                                 8x GE SFP gPIM
    Xcvr 0       REV 01   740-013111   8043353           SFP-T
    Xcvr 1                NON-JNPR     PC602QW           SFP-SX
    Xcvr 2       k        NON-JNPR     BDS3I             SFP-1000BASE-BX10-D
    Xcvr 3       REV 01   740-011612   9XT702501080      SFP-LH
    Xcvr 4       REV 01   740-011612   9XT702501079      SFP-LH
    Xcvr 5                NON-JNPR     PCH2GTJ           SFP-SX
    Xcvr 6                NON-JNPR     PC604DL           SFP-SX
    Xcvr 7       REV 01   740-011620   5349504           SFP-FX
FPC 8            REV 00   750-038290                     FPC
Power Supply 0
```

Meaning The output displays the hardware details of the device and a list of all interfaces configured.

Verify that the output contains the following values:

- **FPC 5, PIC 0** —8x SFP gPIM
- **FPC 6, PIC 0** —8x SFP gPIM

NOTE: In the example, the output for 8-Port SFP Gigabit Ethernet XPIM is displayed as 8x GE SFP gPIM.

Verifying the FPC Status

Purpose Verify that the status of the Flexible PIC Concentrator is online.

Action From operational mode, enter the **show chassis fpc pic-status** command.

```
user@host>  show chassis fpc pic-status
Slot 0   Online       FPC
  PIC 0  Online       4x GE Base PIC
Slot 1   Present      FPC
Slot 5   Online       FPC
  PIC 0  Online       8x GE SFP gPIM
Slot 6   Online       FPC
  PIC 0  Online       8x GE SFP gPIM
Slot 8   Present      FPC
```

Meaning The output shows the FPC status for slot 5 and slot 6 as online. The 8-Port Gigabit Ethernet SFP XPIM is installed in slot 5 and slot 6 of the device.

Verifying Interface Link Status on Device 1

Purpose Verify that the interface link status is up.

Action From operational mode, enter the **show interface terse ge-6/0/*** command.

```
user@host> show interface terse ge-6/0/*
```

Output for Device 1

```
Interface               Admin Link Proto    Local                 Remote
ge-6/0/0                up    up
ge-6/0/0.0              up    up   inet     10.1.1.1/24
ge-6/0/1                up    up
ge-6/0/1.0              up    up   inet     11.1.1.1/24
ge-6/0/2                up    up
ge-6/0/2.0              up    up   inet     12.1.1.1/24
ge-6/0/3                up    up
ge-6/0/3.0              up    up   inet     13.1.1.1/24
ge-6/0/4                up    up
ge-6/0/4.0              up    up   inet     14.1.1.1/24
ge-6/0/5                up    up
ge-6/0/5.0              up    up   inet     15.1.1.1/24
ge-6/0/6                up    up
ge-6/0/6.0              up    up   inet     16.1.1.1/24
ge-6/0/7                up    up
ge-6/0/7.0              up    up   inet     17.1.1.1/24
```

Meaning The output displays a list of all interfaces configured.

If the link displays **up** for all interfaces, the configuration is working properly. This verifies that the XPIM is up and end-to-end ping is working.

Verifying the Interface Settings on Device 1

Purpose Verify that the interfaces are configured as expected.

Action From operational mode, enter the **show interface ge-6/0/0 extensive | no-more** command.

```
user@host>show interface ge-6/0/0 extensive | no-more
```

Output for Device 1

```
Physical interface: ge-6/0/0, Enabled, Physical link is Up
  Interface index: 152, SNMP ifIndex: 544, Generation: 155
  Link-level type: Ethernet, MTU: 9192, 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
  Interface flags: SNMP-Traps Internal: 0x0
```

```
Link flags     : None
CoS queues     : 8 supported, 8 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:26:88:04:0a:a8, Hardware address: 00:26:88:04:0a:a8
Last flapped   : 2012-07-05 21:58:46 PDT (00:13:29 ago)
Statistics last cleared: Never
Traffic statistics:
 Input  bytes  :                  228                    0 bps
 Output bytes  :                  540                    0 bps
 Input  packets:                    3                    0 pps
 Output packets:                    6                    0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets

  0 best-effort                    3                    3                    0

  1 expedited-fo                   0                    0                    0

  2 assured-forw                   0                    0                    0

  3 network-cont                   0                    0                    0

Queue number:         Mapped forwarding classes
  0                   best-effort
  1                   expedited-forwarding
  2                   assured-forwarding
  3                   network-control
Active alarms  : None
Active defects : None
MAC statistics:                       Receive         Transmit
  Total octets                            268              268
  Total packets                             3                3
  Unicast packets                           3                2
  Broadcast packets                         0                1
  Multicast packets                         0                0
  CRC/Align errors                          0                0
  FIFO errors                               0                0
  MAC control frames                        0                0
  MAC pause frames                          0                0
  Oversized frames                          0
  Jabber frames                             0
  Fragment frames                           0
  VLAN tagged frames                        0
  Code violations                           0
Filter statistics:
  Input packet count                        0
  Input packet rejects                      0
  Input DA rejects                          0
  Input SA rejects                          0
  Output packet count                                        0
  Output packet pad count                                    0
  Output packet error count                                  0
  CAM destination filters: 2, CAM source filters: 0
```

```
  Autonegotiation information:
    Negotiation status: Complete
    Link partner:
        Link mode: Full-duplex, Flow control: None, Remote fault: OK,
        Link partner Speed: 1000 Mbps
    Local resolution:
        Flow control: None, Remote fault: Link OK
  Packet Forwarding Engine configuration:
    Destination slot: 6
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority
Limit
                              %           bps     %           usec
    0 best-effort            95     950000000    95              0      low
none
    3 network-control         5      50000000     5              0      low
none
  Interface transmit statistics: Disabled

  Logical interface ge-6/0/0.0 (Index 81) (SNMP ifIndex 509) (Generation 146)
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                  42
     Input  packets:                   0
     Output packets:                   1
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                  42
     Input  packets:                   0
     Output packets:                   1
    Transit statistics:
     Input  bytes  :                   0                  0 bps
     Output bytes  :                   0                  0 bps
     Input  packets:                   0                  0 pps
     Output packets:                   0                  0 pps
    Security: Zone: HOST
    Allowed host-inbound traffic : any-service bfd bgp dvmrp igmp ldp msdp nhrp
    ospf ospf3 pgm pim rip ripng router-discovery rsvp sap vrrp
    Flow Statistics :
    Flow Input statistics :
      Self packets :                  0
      ICMP packets :                  0
      VPN packets :                   0
      Multicast packets :             0
      Bytes permitted by policy :     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: 9178, Generation: 162, Route table: 0
    Flags: Sendbcast-pkt-to-re
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: 10.1.1.255,
      Generation: 176
```

Meaning The output displays a list of all interface verification parameters.

Verify the following information in the command output:

- Physical Interface—ge-6/0/0, enabled, physical link is **Up**
- MTU—9192
- Speed—1000 Mbps

If the verification parameters are as expected, the configuration is working properly.

Verifying Interface Link Status on Device 2

Purpose Verify that the interface link status is up.

Action From operational mode, enter the **show interface terse ge-6/0/*** command.

```
user@host> show interface terse ge-6/0/*
```

Output for Device 2

```
Interface               Admin Link Proto    Local                 Remote
ge-6/0/0                up    up
ge-6/0/0.0              up    up   inet     10.1.1.2/24
ge-6/0/1                up    up
ge-6/0/1.0              up    up   inet     11.1.1.2/24
ge-6/0/2                up    up
ge-6/0/2.0              up    up   inet     12.1.1.2/24
ge-6/0/3                up    up
ge-6/0/3.0              up    up   inet     13.1.1.2/24
ge-6/0/4                up    up
ge-6/0/4.0              up    up   inet     14.1.1.2/24
ge-6/0/5                up    up
ge-6/0/5.0              up    up   inet     15.1.1.2/24
ge-6/0/6                up    up
ge-6/0/6.0              up    up   inet     16.1.1.2/24
ge-6/0/7                up    up
ge-6/0/7.0              up    up   inet     17.1.1.2/24
```

Meaning The output displays a list of all interfaces configured.

If the link displays **up** for all interfaces, the configuration is working properly. This verifies that the XPIM is up and end-to-end ping is working.

Verifying the Interface Settings on Device 2

Purpose Verify that the interfaces are configured as expected.

Action From operational mode, enter the **show interface ge-6/0/0 extensive | no-more** command.

```
user@host>show interface ge-6/0/0 extensive | no-more
```

Output for Device 2

```
Physical interface: ge-6/0/0, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 520, Generation: 147
  Link-level type: Ethernet, MTU: 9192, 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
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:24:dc:17:2f:a8, Hardware address: 00:24:dc:17:2f:a8
  Last flapped   : 2012-07-05 21:59:42 PDT (00:15:32 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                 228                 0 bps
   Output bytes  :                 294                 0 bps
   Input  packets:                   3                 0 pps
   Output packets:                   5                 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: 13, 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                   3                   3                   0

    1 expedited-fo                  0                   0                   0

    2 assured-forw                  0                   0                   0

    3 network-cont                  0                   0                   0

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

```
  Active alarms  : None
  Active defects : None
  MAC statistics:                      Receive         Transmit
    Total octets                           268              268
    Total packets                            3                3
    Unicast packets                          2                3
    Broadcast packets                        1                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: Complete
    Link partner:
        Link mode: Full-duplex, Flow control: None, Remote fault: OK,
        Link partner Speed: 1000 Mbps
    Local resolution:
        Flow control: None, Remote fault: Link OK
  Packet Forwarding Engine configuration:
    Destination slot: 6
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority
Limit
                                 %           bps     %           usec
    0 best-effort               95     950000000    95              0      low
none
    3 network-control            5      50000000     5              0      low
none
  Interface transmit statistics: Disabled

  Logical interface ge-6/0/0.0 (Index 73) (SNMP ifIndex 509) (Generation 146)
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                   42
     Input  packets:                    0
     Output packets:                    1
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                   42
     Input  packets:                    0
     Output packets:                    1
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
```

```
 Input  packets:                    0                    0 pps
 Output packets:                    0                    0 pps
Security: Zone: HOST
Allowed host-inbound traffic : any-service bfd bgp dvmrp igmp ldp msdp nhrp
ospf ospf3 pgm pim rip ripng router-discovery rsvp sap vrrp
Flow Statistics :
Flow Input statistics :
  Self packets :                   0
  ICMP packets :                   0
  VPN packets :                    0
  Multicast packets :              0
  Bytes permitted by policy :      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: 9178, Generation: 162, Route table: 0
  Flags: Sendbcast-pkt-to-re
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.1.1/24, Local: 10.1.1.2, Broadcast: 10.1.1.255,
    Generation: 176
```

Meaning The output displays a list of all interface verification parameters.

Verify the following information in the command output:

- Physical Interface—ge-6/0/0, enabled, physical link is **Up**
- MTU—9192
- Speed—1000 Mbps

If the verification parameters are as expected, the configuration is working properly.

Related Documentation

- Understanding the 8-Port Gigabit Ethernet SFP XPIM on page 25
- *Ethernet Interfaces Feature Guide for Security Devices*

Example: Configuring the 2-Port 10-Gigabit Ethernet XPIM Interface

Supported Platforms LN Series, SRX650

This example shows how to perform basic configuration for the 1-Port Gigabit Ethernet SFP Mini-PIM.

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 51.

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 commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set interfaces xe-6/0/0 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

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.

```
   Temp   CPU Utilization  (%)  Memory   Utilization (%)
Slot 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.

NOTE: The 2-Port 10-Gigabit Ethernet XPIM is installed in the second and sixth slot of the SRX650 device chassis; therefore the FPC used here is fpc 2 and fpc 6.

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: Sendbcast-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 27
- Understanding Ethernet Interfaces on page 3
- *Ethernet Interfaces Feature Guide for Security Devices*

CHAPTER 12

Ethernet OAM Link Fault Management

Example: Configuring Ethernet OAM Link Fault Management

Supported Platforms LN Series, SRX100, SRX210, SRX220, SRX240, SRX550, SRX650

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 feature is supported on SRX100, SRX210, SRX220, SRX240, SRX550, and SRX650 devices.

This example describes how to enable and configure OAM LFM on a Gigabit Ethernet or Fast Ethernet interface:

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 51.
- 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 31

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 31.

Figure 3 on page 96 shows the topology used in this example.

Figure 3: Ethernet LFM with SRX Series Devices

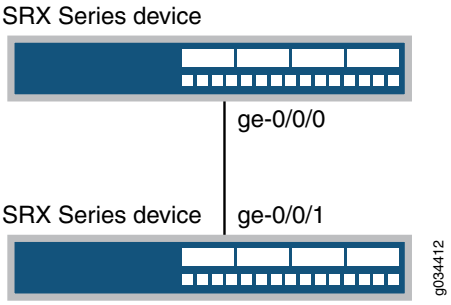

NOTE: For more information about configuring Ethernet OAM Link Fault Management, see Junos® OS Ethernet Interfaces.

Configuration

To configure Ethernet OAM LFM, perform these tasks:

- Configuring Ethernet OAM Link Fault Management on Device 1 on page 96
- Configuring Ethernet OAM Link Fault Management on Device 2 on page 97

Configuring Ethernet OAM Link Fault Management on Device 1

CLI Quick Configuration

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

```
set 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*.

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 commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set 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

- *Ethernet Port Switching Feature Guide for Security Devices*
- Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways on page 31
- *Ethernet Interfaces Feature Guide for Security Devices*

CHAPTER 13

Power over Ethernet

Example: Configuring PoE on All Interfaces

Supported Platforms LN Series, SRX210, SRX220, SRX240, SRX650

This example shows how to configure PoE on all interfaces.

Requirements

Before you begin, configure Ethernet interfaces. See "Example: Creating an Ethernet Interface" on page 51.

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 commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set 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 *Using the CLI Editor in Configuration Mode*.

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 SRX240 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 35
- Example: Configuring PoE on an Individual Interface on page 103
- Example: Disabling a PoE Interface on page 106
- *Ethernet Interfaces Feature Guide for Security Devices*

Example: Configuring PoE on an Individual Interface

Supported Platforms LN Series, SRX210, SRX220, SRX240, SRX650

This example shows how to configure PoE on an individual interface.

Requirements

Before you begin:

- Configure Ethernet interfaces. See "Example: Creating an Ethernet Interface" on page 51.
- Configure PoE on all interfaces. See "Example: Configuring PoE on All Interfaces" on page 101.

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 commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set 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 *Using the CLI Editor in Configuration Mode*.

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
   ```

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 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

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 SRX240 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 SRX240 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
Sl 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-0/0/1 5
Sl 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        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 35
- Example: Configuring PoE on All Interfaces on page 101
- Example: Disabling a PoE Interface on page 106
- *Ethernet Interfaces Feature Guide for Security Devices*

Example: Disabling a PoE Interface

Supported Platforms LN Series, SRX210, SRX220, SRX240, SRX650

This example shows how to disable PoE on all interfaces or on a specific interface.

Requirements

Before you begin:

- Configure PoE on all interfaces. See "Example: Configuring PoE on All Interfaces" on page 101.
- Configure PoE on an individual interface. See "Example: Configuring PoE on an Individual Interface" on page 103.

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 35
- *Ethernet Interfaces Feature Guide for Security Devices*

CHAPTER 14

Ethernet Interface Configuration Statements

Interfaces Configuration Statement Hierarchy

Supported Platforms J Series, LN Series, SRX Series

Use the statements in the **interfaces** configuration hierarchy to configure interfaces on the device.

```
interfaces {
    interface-name {
        accounting-profile name;
        clocking (external | internal);
        dce;
        description text;
        disable;
        e1-options {
            bert-algorithm algorithm;
            bert-error-rate rate;
            bert-period seconds;
            fcs (16 | 32);
            framing (g704 | g704-no-crc4 | unframed);
            idle-cycle-flag (flags | ones);
            invert-data data;
            loopback (local | remote);
            start-end-flag (shared | filler);
            timeslots time-slot-range;
        }
        e3-options {
            bert-algorithm algorithm;
            bert-error-rate rate;
            bert-period seconds;
            compatibility-mode {
                digital-link {
                    subrate value;
                }
                kentrox {
                    subrate value;
                }
                larscom;
            }
            fcs (16 | 32);
            framing (g.751 | g.832);
            idle-cycle-flag value;
            invert-data;
            loopback (local | remote);
            (no-payload-scrambler | payload-scrambler);
            (no-unframed | -unframed);
            start-end-flag (filler | shared);
        }
        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);
        fastether-options {
            802.3ad interface-name {
                (backup | primary);
                lacp {
                    port-priority port-number;
                }
            }
```

```
    (auto-negotiation | no-auto-negotiation);
    ignore-l3-incompletes;
    ingress-rate-limit rate;
    (loopback | no-loopback);
    mpls {
      pop-all-labels {
        required-depth number;
      }
    }
    redundant-parent interface-name;
    source-address-filter mac-address;
  }
  flexible-vlan-tagging;
  gigether-options {
    802.3ad interface-name {
      (backup | primary);
      lacp {
        port-priority port-number;
      }
    }
    (auto-negotiation <remote-fault> (local-interface-offline | local-interface-online)
      | no-auto-negotiation);
    (flow-control | no-flow-control);
    ignore-l3-incompletes;
    (loopback | no-loopback);
    mpls {
      pop-all-labels {
        required-depth [number];
      }
    }
    redundant-parent interface-name;
    source-address-filter mac-address;
  }
  gratuitous-arp-reply;
  hierarchical-scheduler {
    maximum-hierarchy-levels 2;
  }
  hold-time {
    down milliseconds;
    up milliseconds;
  }
  keepalives {
    down-count number;
    interval number;
    up-count number;
  }
  link-mode (full-duplex | half-duplex);
  lmi {
    lmi-type (ansi | c-lmi | itu);
    n391dte number;
    n392dce number;
    n392dte number;
    n393dce number;
    n393dte number;
    t391dte number;
    t392dce number;
```

```
}
logical-tunnel-options {
  per-unit-mac-disable;
}
mac mac-address;
mtu bytes;
native-vlan-idvlan-id;
no-gratuitous-arp-request;
no-keepalives;
optics-options {
  alarm {
    low-light-alarm (link-down | syslog);
  }
  warning {
    low-light-warning (link-down | syslog);
  }
  wavelength wavelength-options;
}
otn-options {
  bytes {
    transmit-payload-type number];
  }
  fec (efec | gfec | none);
  (laser-enable | no-laser-enable);
  (line-loopback | no-line-loopback);
  rate (fixed-stuff-bytes | no-fixed-stuff-bytes | pass-thru);
  trigger {
    oc-lof {
      hold-time {
        down milliseconds;
        up milliseconds;
        }
      ignore;
    }
    oc-lom {
      hold-time {
        down milliseconds;
        up milliseconds;
        }
      ignore;
    }
    oc-los {
      hold-time {
        down milliseconds;
        up milliseconds;
        }
      ignore;
    }
    oc-wavelength-lock {
      hold-time {
        down milliseconds;
        up milliseconds;
        }
      ignore;
    }
    odu-ais {
```

```
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  odu-bdi {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  odu-lck {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  odu-oci {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  odu-sd {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  odu-tca-bbe {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  odu-tca-es {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  odu-tca-ses {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  odu-tca-uas {
```

```
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  odu-ttim {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  opu-ptim {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-ais {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-bdi {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-bdi {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-fec-deg {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-fec-deg {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-fec-exe {
```

```
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-iae {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-sd {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-tca-bbe {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-tca-es {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-tca-ses {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-tca-uas {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
  otu-ttim {
    hold-time {
      down milliseconds;
      up milliseconds;
      }
    ignore;
  }
}
```

```
    tti (odu-dapi | odu-expected-receive-dapi | odu-expected-receive-sapi | odu-sapi |
      otu-dapi |otu-expected-receive-dapi | otu-expected-receive-sapi |otu-sapi);
  }
  passive-monitor-mode;
  (per-unit-scheduler | no-per-unit-schedule);
  port-mirror-instance;
  ppp-options {
    chap {
      access-profile name;;
      default-chap-secret secret;
      local-name name;
      no-rfc2486;
      passive;
    }
    compression {
      acfc;
      pfc;
    }
    dynamic-profile (dynamic-profile | junos-default-profile);
    lcp-max-conf-req number;
    lcp-restart-timer milliseconds;
    loopback-clear-timer seconds;
    ncp-max-conf-req number;
    ncp-restart-timer milliseconds;
    no-termination-request;
    pap {
      access-profile name;
      default-password password;
      local-name name;
      local-password password;
      no-rfc2486;
      passive;
    }
  }
  promiscuous-mode;
  receive-bucket {
    overflow {
      discard;
      tag;
    }
    rate number;
    threshold number;
  }
  redundant-pseudo-interface-options {
    redundancy-group number;
  }
  satop-options {
    excessive-packet-loss-rate {
      sample-period milliseconds;
      threshold percentage;
    }
    idle-pattern number;
    (jitter-buffer-auto-adjust | jitter-buffer-latency milliseconds | jitter-buffer-packets
      number;
    payload-size number;
  }
```

```
speed (100m |10m | 1g);
stacked-vlan-tagging;
switch-options {
  switch-port port-number {
    (auto-negotiation | no-auto-negotiation);
    cascade-port;
    link-mode (full-duplex | half-duplex);
    speed (100m |10m | 1g);
    vlan-id number;
  }
}
t1-options {
  alarm-compliance {
    accunet-t1-5-service;
  }
  bert-algorithm algorithm;
  bert-error-rate rate;
  bert-period seconds;
  buildout value;
  byte-encoding (nx56 | nx64);
  fcs (16 | 32);
  framing (esf | sf);
  idle-cycle-flags (flags | ones);
  invert-data;
  line-encoding (ami | b8zs);
  loopback (local | payload | remote);
  remote-loopback-respond;
  start-end-flag (filler | shared);
  timeslots time-slot-range;
}
t3-options {
  bert-algorithm algorithm ;
  bert-error-rate rate ;
  bert-period seconds ;
  (cbit-parity | no-cbit-parity);
  compatibility-mode {
    adtran {
      subrate value;
    }
    digital-link {
      subrate value;
    }
    kentrox {
      subrate value;
    }
    larscom;
      subrate value;
    }
    verilink;
      subrate value;
    }
  }
  fcs (16 | 32);
  (feac-loop-respond | no-feac-loop-respond);
  idle-cycle-flag (flags | ones);
  (long-buildout | no-long-buildout);
```

```
    (loop-timing | no-loop-timing);
    loopback (local | payload | remote);
    (no-payload-scrambler | payload-scrambler);
    (no-unframed | unframed);
    start-end-flag value (filler | shared);
}
traceoptions {
    flag (all | event | ipc | media);
}
transmit-bucket {
    overflow {
        discard;
    }
    rate number;
    threshold number;
}
(traps | no-traps);
unit unit-number {
    accept-source-mac {
        mac-address mac-address;
    }
    accounting-profile name;
    arp-resp (restricted | unrestricted);
    backup-options {
        interface interface-name;
    }
    bandwidth bandwidth;
    description text;
    disable;
    encapsulation (dix | ether-vpls-fr | frame-relay-ppp | ppp-over-ether | vlan-bridge |
        vlan-ccc | vlan-vpls |vlan-tcc);
    family {
        bridge {
            bridge-domain-type (svlan| bvlan);
            filter {
                group number;
                input filter-name;
                input-list [filter-name];
                output filter-name;
                output-list [filter-name];
            }
            interface-mode (access | trunk);
            policer {
                input input-policer-name;
                output outputpolicer-name;
            }
            vlan-id vlan-id;
            vlan-id-list [vlan-id];
            vlan-rewrite {
                translate {
                    from-vlan-id;
                    to-vlan-id ;
                }
            }
        }
        ccc {
```

```
    filter {
      group number;
      input filter-name;
      input-list [filter-name];
      output filter-name;
      output-list [filter-name];
    }
    policer {
      input input-policer-name;
      output output-policer-name;
    }
  }
  ethernet-switching {
    native-vlan-id native-vlan-id;
    port-mode (access | tagged-access | trunk);
    reflective-relay;
    vlan {
      members [member-name];
    }
  }
  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-timeseconds> | 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;
    }
  }
  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 | statefull);
          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;
        }
      }
```

```
    iso {
        address source-address;
        mtu value;
    }
    mlfr-end-to-end {
        bundle bundle-name;
    }
    mlfr-uni-nni {
        bundle bundle-name;
    }
    mlppp {
        bundle bundle-name;
    }
    mpls {
        filter {
            group number;
            input filter-name;
            input-list [filter-name];
            output filter-name;
            output-list [filter-name];
        }
        mtu mtu-value;
        policer {
            input input-name;
            output output-name;
        }
    }
    tcc {
        policer {
            input input-name;
            output output-name;
        }
        proxy {
            inet-address inet-address;
        }
        remote {
            inet-address inet-address;
            mac-address mac-address;
        }
    }
    vpls {
        filter {
            group number;
            input filter-name;
            input-list [filter-name];
            output filter-name;
            output-list [filter-name];
        }
        policer {
            input input-name;
            output output-name;
        }
    }
}
input-vlan-map {
    inner-tag-protocol-id tpid;
```

```
    inner-vlan-id number ;
    (pop | push | swap);
    tag-protocol-id tpid;
    vlan-id number;
  }
  interface-shared-with {
    psd-name;
  }
  native-inner-vlan-id value;
  (no-traps | traps);
  output-vlan-map {
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    (pop | push | swap);
    tag-protocol-id tpid;
    vlan-id number;
  }
  ppp-options {
    chap {
      access-profile name;
      default-chap-secret name;
      local-name name;
      no-rfc2486;
      passive;
    }
    dynamic-profile profile-name;
    lcp-max-conf-req number;
    lcp-restart-timer milliseconds;
    loopback-clear-timer seconds;
    ncp-max-conf-req number;
    ncp-restart-timer milliseconds;
    no-termination-request;
    pap {
      access-profile name;
      default-password password;
      local-name name;
      local-password password;
      no-rfc2486;
      passive;
    }
  }
  proxy-arp (restricted | unrestricted);
  radio-router {
    bandwidth number;
    credit {
      interval number;
    }
    data-rate number;
    latency number;
    quality number;
    resource number;
    threshold number;
  }
  swap-by-poppush;
  traps;
  vlan-id vlan-id;
```

```
        vlan-id-range vlan-id-range;
        vlan-id-list [vlan-id];
        vlan-id-range vlan-id1-vlan-id2;
        vlan-tags {
          (inner vlan-id | inner-range vlan-id1-vlan-id2);
          inner-list [vlan-id];
          outer vlan-id;
        }
      }
      vlan-tagging;
    }
  }
```

Related Documentation

- *Layer 2 Bridging and Transparent Mode Feature Guide for Security Devices*
- *Junos OS Interfaces Library for Security Devices*
- *Administration Guide for Security Devices*

Chassis Configuration Statement Hierarchy

Supported Platforms J Series, LN Series, SRX Series

Use the statements in the **chassis** configuration hierarchy to configure alarms, aggregated devices, clusters, the routing engine, and other chassis properties.

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

```
    serial {
      cts-absent (ignore | red | yellow);
      dcd-absent (ignore | red | yellow);
      dsr-absent (ignore | red | yellow);
      loss-of-rx-clock (ignore | red | yellow);
      loss-of-tx-clock (ignore | red | yellow);
    }
    services {
      hw-down (ignore | red | yellow);
      linkdown (ignore | red | yellow);
      pic-hold-reset (ignore | red | yellow);
      pic-reset (ignore | red | yellow);
      rx-errors (ignore | red | yellow);
      sw-down (ignore | red | yellow);
      tx-errors (ignore | red | yellow);
    }
    t3 {
      ais (ignore | red | yellow);
      exz (ignore | red | yellow);
      ferf (ignore | red | yellow);
      idle (ignore | red | yellow);
      lcv (ignore | red | yellow);
      lof (ignore | red | yellow);
      los (ignore | red | yellow);
      pll (ignore | red | yellow);
      ylw (ignore | red | yellow);
    }
  }
  cluster {
    control-link-recovery;
    heartbeat-interval milliseconds;
    heartbeat-threshold number;
    network-management {
      cluster-master;
    }
    redundancy-group group-number {
      gratuitous-arp-count number;
      hold-down-interval number;
      interface-monitor interface-name {
        weight number;
      }
      ip-monitoring {
        family {
          inet {
            ipv4-address {
              interface {
                logical-interface-name;
                secondary-ip-address ip-address;
              }
              weight number;
            }
          }
        }
        global-threshold number;
        global-weight number;
        retry-count number;
```

```
            retry-interval seconds;
        }
        node (0 | 1 ) {
            priority number;
        }
        preempt;
    }
    reth-count number;
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            (world-readable | no-world-readable);
            size maximum-file-size;
        }
        flag flag;
        level {
            (alert | all | critical | debug | emergency | error | info | notice | warning);
        }
        no-remote-trace;
    }
}
config-button {
    no-clear;
    no-rescue;
}
craft-lockout;
fpc slot-number {
    offline;
    pic slot-number {
        aggregate-ports;
        framing {
            (e1 | e3 | sdh | sonet | t1 | t3);
        }
        max-queues-per-interface (4 | 8);
        mlfr-uni-nni-bundles number;
        no-multi-rate;
        port slot-number {
            framing (e1 | e3 | sdh | sonet | t1 | t3);
            speed (oc12-stm4 | oc3-stm1 | oc48-stm16);
        }
        q-pic-large-buffer (large-scale | small-scale);
        services-offload {
            low-latency;
            per-session-statistics;
        }
        shdsl {
            pic-mode (1-port-atm | 2-port-atm | 4-port-atm | efm);
        }
        sparse-dlcis;
        traffic-manager {
            egress-shaping-overhead number;
            ingress-shaping-overhead number;
            mode (egress-only | ingress-and-egress);
        }
```

```
        tunnel-queuing;
      }
    }
    ioc-npc-connectivity {
      ioc slot-number {
        npc (npc-slot-number | none);
      }
    }
    maximum-ecmp (16 | 32 | 64);
    network-services (ethernet | IP);
    routing-engine {
      bios {
        no-auto-upgrade;
      }
      on-disk-failure {
        disk-failure-action (halt | reboot);
      }
      usb-wwan {
        port 1;
      }
    }
    usb {
      storage {
        disable;
      }
    }
  }
```

Related Documentation

- *IP Monitoring Feature Guide for Security Devices*
- *Chassis Cluster Feature Guide for Security Devices*
- *Master Administrator for Logical Systems Feature Guide for Security Devices*

encapsulation (Interfaces)

Supported Platforms J Series, LN Series, SRX Series

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 *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Release 9.5 of Junos OS.

Description Specify logical link layer encapsulation.

Options

- **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

- *Junos OS Interfaces Library for Security Devices*

family inet (Interfaces)

Supported Platforms J Series, LN Series, SRX Series

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-timeseconds> | 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

- *Ethernet Port Switching Feature Guide for Security Devices*
- *Layer 2 Bridging and Transparent Mode Feature Guide for Security Devices*
- *Junos OS Interfaces Library for Security Devices*

family inet6

Supported Platforms J Series, LN Series, SRX Series

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 | statefull);
    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 and J 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
- *Layer 2 Bridging and Transparent Mode Feature Guide for Security Devices*
- *Junos OS Interfaces Library for Security Devices*

flow-control (Interfaces)

Supported Platforms J Series, LN Series, SRX Series

Syntax
```
(flow-control | no-flow-control);
```

Hierarchy Level
```
[edit interfaces interface-name fastether-options],
[edit interfaces interface-name gigether-options],
[edit interfaces interface-name redundant-ether-options]
```

Release Information Statement modified in Release 9.0 of Junos OS.

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.

NOTE: The flow-control option is not supported on Fast Ethernet and Gigabit Ethernet interfaces on J Series devices.

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
- *Junos OS Interfaces Library for Security Devices*

link-speed (Interfaces)

Supported Platforms J Series, LN Series, SRX Series

Syntax **link-speed** ***speed*;**

Hierarchy Level **[edit interfaces** ***interface-name*** **redundant-ether-options]**

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 ***speed*** —For redundant Ethernet links, you can specify ***speed*** in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

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

Related Documentation
- *Junos OS Interfaces Library for Security Devices*

lacp (Interfaces)

Supported Platforms J Series, LN Series, SRX Series

Syntax
```
lacp {
  port-priority port-number;
}
```

Hierarchy Level **[edit interfaces** ***interface-name*** **redundant-ether-options]**

Release Information Statement introduced in Release 10.2 of Junos OS.

Description For redundant Ethernet interfaces in a chassis cluster only, configure Link Aggregation Control Protocol (LACP).

Options
- **active**—Initiate transmission of LACP packets.
- **passive**—Respond to LACP packets.

Default: If you do not specify **lacp** as either **active** or **passive**, LACP remains off (the default).

The remaining statements are explained separately.

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

Related Documentation
- *Junos OS Interfaces Library for Security Devices*

loopback (Interfaces)

Supported Platforms J Series, LN Series, SRX Series

Syntax **(loopback | no-loopback);**

Hierarchy Level **[edit interfaces *interface-name* redundant-ether-options]**

Release Information Statement modified in Release 9.0 of Junos OS.

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
- *Junos OS Interfaces Library for Security Devices*

media-type (Interfaces)

Supported Platforms LN Series, SRX550, SRX650

Syntax media-type

Hierarchy Level [edit interfaces *interface-name* media-type]

Release Information Command introduced in Release 10.2 of Junos OS.

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
- *Junos OS Interfaces Library for Security Devices*

minimum-links (Interfaces)

Supported Platforms J Series, LN Series, SRX Series

Syntax **minimum-links** ***number***;

Hierarchy Level **[edit interfaces** ***interface-name*** **redundant-ether-options]**

Release Information Statement added in Release 10.1 of Junos OS.

Description 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.

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.

Options ***number***—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.

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

Related Documentation

- *Junos OS Interfaces Library for Security Devices*

periodic (Interfaces)

Supported Platforms J Series, LN Series, SRX Series

Syntax **periodic (fast | slow);**

Hierarchy Level **[edit interfaces *interface-name* redundant-ether-options lacp]**

Release Information Statement introduced in Release 10.2 of Junos OS.

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

- **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

- *Junos OS Interfaces Library for Security Devices*

ppp-over-ether

Supported Platforms J Series, LN Series, SRX100, SRX110, SRX210, SRX220, SRX240, SRX550, SRX650

Syntax ppp-over-ether;

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

Release Information Statement introduced before Release 11.2 of Junos OS.
This encapsulation is supported for Redundant Ethernet interface in Release 11.2 of Junos OS.

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

- *Junos OS Interfaces Library for Security Devices*

promiscuous-mode (Interfaces)

Supported Platforms J Series, LN Series, SRX Series

Syntax
```
promiscuous-mode;
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced in Release 10.1 of Junos OS.

Description 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.

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.

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

Related Documentation
- *Junos OS Interfaces Library for Security Devices*

r2cp

Supported Platforms LN Series, SRX100, SRX110, SRX210, SRX220, SRX240, SRX550, SRX650

Syntax
```
r2cp {
  command binary-file-path;
  disable;
}
```

Hierarchy Level [edit system processes]

Release Information Statement introduced in Junos OS Release 8.5.

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
- **command *binary-file-path***—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
- *Junos OS Interfaces Library for Security Devices*

redundancy-group (Interfaces)

Supported Platforms J Series, LN Series, SRX Series

Syntax redundancy-group *number*;

Hierarchy Level **[edit interfaces *interface-name* redundant-ether-options]**

Release Information Statement introduced in Release 9.0 of Junos OS.

Description Specify the redundancy group that a redundant Ethernet interface belongs to.

Options ***number*** —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

- *Ethernet Port Switching Feature Guide for Security Devices*
- *Layer 2 Bridging and Transparent Mode Feature Guide for Security Devices*
- *Junos OS Interfaces Library for Security Devices*

redundant-ether-options

Supported Platforms J Series, LN Series, SRX Series

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 Release 9.0 of Junos OS.

Description Configure Ethernet redundancy options for a chassis cluster.

Options The remaining statements are explained separately.

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

Related Documentation
- *Junos OS Interfaces Library for Security Devices*

redundant-parent (Interfaces Gigabit Ethernet)

Supported Platforms J Series, LN Series, SRX Series

Syntax redundant-parent *interface-name*;

Hierarchy Level [edit interfaces *interface-name* gigether-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 ***interface*** —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
- *Junos OS Interfaces Library for Security Devices*

redundant-parent (Interfaces Fast Ethernet)

Supported Platforms J Series, LN Series, SRX Series

Syntax **redundant-parent** ***interface-name*** **;**

Hierarchy Level **[edit interfaces** ***interface-name*** **fastether-options]**

Release Information Statement introduced in Release 9.0 of Junos OS.

Description Configure Fast Ethernet-specific interface properties for Ethernet redundancy in a chassis cluster.

Options ***interface*** —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

- *Junos OS Interfaces Library for Security Devices*

source-address-filter (Interfaces)

Supported Platforms LN Series, SRX100, SRX210, SRX220, SRX240, SRX550, SRX650

Syntax **source-address-filter** *mac-address* ;

Hierarchy Level **[edit interfaces** *interface-name* **redundant-ether-options]**

Release Information Statement modified in Release 9.0 of Junos OS.

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:

- The source address filtering is not supported on the Gigabit Ethernet Interfaces on J Series devices.
- Software based MAC limiting is supported on SRX100, SRX210, SRX220, SRX240, and SRX650 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:nn*** (for example, **00:11:22:33:44:55**) or ***nnnn:nnnn:nnnn*** (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

- *Junos OS Interfaces Library for Security Devices*

source-filtering (Interfaces)

Supported Platforms LN Series, SRX100, SRX210, SRX220, SRX240, SRX550, SRX650

Syntax **(source-filtering | no-source-filtering);**

Hierarchy Level **[edit interfaces *interface-name* redundant-ether-options]**

Release Information Statement modified in Release 9.0 of Junos OS.

Description 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.

If the remote Ethernet card is changed, the interface cannot receive packets from the new card because it has a different MAC address.

By default, source address filtering is disabled.

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

Related Documentation
- *Junos OS Interfaces Library for Security Devices*

speed (Interfaces)

Supported Platforms LN Series, SRX550, SRX650

Syntax speed (100m |10m | 1g);

Hierarchy Level [edit interfaces *interface-name* speed]

Release Information Command introduced in Release 10.2 of Junos OS.

Description Configure the operating speed for the 2-Port 10 Gigabit Ethernet XPIM.

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

Related Documentation
- *Junos OS Interfaces Library for Security Devices*

vlan-tagging (Interfaces)

Supported Platforms LN Series, SRX Series

Syntax vlan-tagging native-vlan-id *vlan-id*;

Hierarchy Level [edit interfaces *interface*]

Release Information Statement introduced in Release 9.5 of Junos OS.

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

Related Documentation

- *Ethernet Port Switching Feature Guide for Security Devices*
- *Junos OS Interfaces Library for Security Devices*
- *Layer 2 Bridging and Transparent Mode Feature Guide for Security Devices*
- *Junos OS Layer 2 Bridging and Switching Library for Security Devices*

CHAPTER 15

PoE Configuration Statements

PoE Configuration Statement Hierarchy

Supported Platforms LN Series, SRX210, SRX220, SRX240, SRX5400, SRX550, SRX650

To configure Power over Ethernet options, use the configuration statements in the **poe** configuration hierarchy. Statement descriptions that are exclusive to the J Series and SRX Series devices running Junos OS are described in this chapter.

```
poe {
   fpc slot-number {
      maximum-power watts;
      priority (high | low);
   }
   guard-band watts;
   interface (all | interface-name) {
      disable;
      maximum-power watts;
      priority (high | low);
      telemetries {
         disable;
         duration hours;
         interval minutes;
      }
   }
   management (class | static);
}
```

Related Documentation
- *Ethernet Interfaces Feature Guide for Security Devices*
- *Junos OS Interfaces Library for Security Devices*

disable (PoE)

Supported Platforms LN Series, SRX210, SRX220, SRX240, SRX5400, SRX550, SRX650

Syntax
```
disable;
```

Hierarchy Level
```
[edit poe interface (all | interface-name) ]
[edit poe interface (all | interface-name) telemetries]
```

Release Information Statement introduced in Release 9.5 of Junos OS.

Description 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.

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

Related Documentation
- *Ethernet Interfaces Feature Guide for Security Devices*
- *Junos OS Interfaces Library for Security Devices*

duration (PoE)

Supported Platforms LN Series, SRX210, SRX220, SRX240, SRX5400, SRX550, SRX650

Syntax duration *hours*;

Hierarchy Level [edit poe interface (all | *interface-name*) telemetries]

Release Information Statement introduced in Release 9.5 of Junos OS.

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

- *Ethernet Interfaces Feature Guide for Security Devices*
- *Junos OS Interfaces Library for Security Devices*

guard-band (PoE)

Supported Platforms LN Series, SRX210, SRX220, SRX240, SRX5400, SRX550, SRX650

Syntax guard-band *watts*;

Hierarchy Level [edit poe]

Release Information Statement introduced in Release 9.5 of Junos OS.

Description Reserves the specified amount of power for the SRX Series device in case of a spike in PoE consumption.

Options ***watts***—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

- *Ethernet Interfaces Feature Guide for Security Devices*
- *Junos OS Interfaces Library for Security Devices*

interface (PoE)

Supported Platforms LN Series, SRX210, SRX220, SRX240, SRX5400, SRX550, SRX650

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 Release 9.5 of Junos OS.

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.

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

Related Documentation

- *Ethernet Interfaces Feature Guide for Security Devices*
- *Junos OS Interfaces Library for Security Devices*

interval (PoE)

Supported Platforms LN Series, SRX210, SRX220, SRX240, SRX5400, SRX550, SRX650

Syntax interval *minutes*;

Hierarchy Level [edit poe interface (all | *interface-name*) telemetries]

Release Information Statement introduced in Release 9.5 of Junos OS.

Description Modifies the interval for logging telemetries if you are monitoring the per-port power consumption for PoE interfaces.

Options ***minutes***—Interval at which data is logged.
Range: 1 through 30 minutes
Default: 5 minutes

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

Related Documentation
- *Ethernet Interfaces Feature Guide for Security Devices*
- *Junos OS Interfaces Library for Security Devices*

management (PoE)

Supported Platforms LN Series, SRX210, SRX220, SRX240, SRX5400, SRX550, SRX650

Syntax management (class | static);

Hierarchy Level [edit poe]

Release Information Statement introduced in Release 9.5 of Junos OS.

Description Designates how the SRX Series device allocates power to the PoE ports.

Default static

Options
- **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
- *Ethernet Interfaces Feature Guide for Security Devices*
- *Junos OS Interfaces Library for Security Devices*

maximum-power (PoE)

Supported Platforms LN Series, SRX210, SRX220, SRX240, SRX5400, SRX550, SRX650

Syntax maximum-power *watts*;

Hierarchy Level [edit poe interface (all | *interface-name*)]

Release Information Statement introduced in Release 9.5 of Junos OS.

Description Maximum amount of power that can be supplied to the port.

Default 15.4 W

Options **Watts**—The maximum number of watts that can be supplied to the port.

Range —0 through 15.4

Default—15.4 W

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

Related Documentation

- *Ethernet Interfaces Feature Guide for Security Devices*
- *Junos OS Interfaces Library for Security Devices*

priority (PoE)

Supported Platforms LN Series, SRX210, SRX220, SRX240, SRX5400, SRX550, SRX650

Syntax priority (high | low);

Hierarchy Level [edit poe interface (all | *interface-name*)]

Release Information Statement introduced in Release 9.5 of Junos OS.

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:

- **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

- *Ethernet Interfaces Feature Guide for Security Devices*
- *Junos OS Interfaces Library for Security Devices*

telemetries (PoE)

Supported Platforms LN Series, SRX210, SRX220, SRX240, SRX5400, SRX550, SRX650

Syntax

```
telemetries {
  disable;
  duration hours;
  interval minutes;
}
```

Hierarchy Level [edit poe interface (all | *interface-name*)]

Release Information Statement introduced in Release 9.5 of Junos OS.

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

- PoE Configuration Statement Hierarchy on page 147
- *Ethernet Interfaces Feature Guide for Security Devices*
- *Junos OS Interfaces Library for Security Devices*

PART 3

Administration

CHAPTER 16

Aggregated Ethernet Interfaces

Verifying Aggregated Ethernet Interfaces

Supported Platforms J Series, LN Series, SRX Series

Verifying Aggregated Ethernet Interfaces (terse)

Supported Platforms J Series, LN Series, 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 J Series, LN Series, 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 57

CHAPTER 17

Link Aggregation Control Protocol

Verifying LACP on Redundant Ethernet Interfaces

Supported Platforms J Series, LN Series, 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:       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

- Example: Configuring LACP on Chassis Clusters on page 66
- Verifying LACP on Standalone Devices on page 160
- *Ethernet Interfaces Feature Guide for Security Devices*
- *Junos OS Interfaces Library for Security Devices*

Verifying LACP on Standalone Devices

Supported Platforms J Series, LN Series, SRX Series

- Verifying LACP Statistics on page 160
- Verifying LACP Aggregated Ethernet Interfaces on page 161

Verifying LACP Statistics

Supported Platforms J Series, LN Series, 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 J Series, LN Series, 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 ifit is collecting and distributing packets.

Related Documentation

- *Ethernet Interfaces Feature Guide for Security Devices*
- *Junos OS Interfaces Library for Security Devices*
- Example: Configuring LACP on Standalone Devices on page 65
- Verifying LACP on Redundant Ethernet Interfaces on page 159

CHAPTER 18

Operational Commands

clear ethernet-switching statistics mac-learning

Supported Platforms LN Series, SRX100, SRX110, SRX210, SRX220, SRX240, SRX550, SRX650

Syntax clear ethernet-switching statistics mac-learning

Release Information Command introduced in Release 10.1 of Junos OS.

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
- show ethernet-switching mac-learning-log
- *Ethernet Interfaces Feature Guide for Security Devices*
- *Junos OS Interfaces Library for Security Devices*

List of Sample Output

clear ethernet-switching statistics mac-learning on page 164
clear ethernet-switching statistics mac-learning interface interface-name on page 164

Sample Output

This command produces no 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 lacp statistics interfaces

Supported Platforms J Series, LN Series, SRX Series

Syntax **clear lacp statistics interfaces** <*interface-name*>

Release Information Command modified in Release 10.2 of Junos OS.

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 196
- *Ethernet Interfaces Feature Guide for Security Devices*
- *Junos OS Interfaces Library for Security Devices*

Output Fields This command produces no output.

show chassis fpc (View)

Supported Platforms LN Series, 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 Release 9.2 of Junos OS; **node**options added in Release 9.0 of Junos OS.

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
- *Flow-Based Processing Feature Guide for Security Devices*

List of Sample Output

Output Fields Table 11 on page 167 lists the output fields for the **show chassis fpc** command. Output fields are listed in the approximate order in which they appear.

Table 11: 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: • **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.
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
                     Temp   CPU Utilization (%)     Memory    Utilization (%)
  Slot State          (C)    Total  Interrupt        DRAM (MB) Heap  Buffer
   0  Online          -------------------- CPU less FPC --------------------
   1  Online          ---------------------  Not Usable --------------------
   2  Online          -------------------- CPU less FPC --------------------
```

Sample Output

show chassis fpc (SRX1400 devices)

```
user@host> show chassis fpc
                     Temp   CPU Utilization (%)      Memory    Utilization (%)
 Slot State           (C)   Total  Interrupt     DRAM (MB)    Heap     Buffer
  0  Online           49     30          0       1024         3          25
  1  Online           41     30          0       1024         3          25
  2  Online           44     30          0       1024         3          25
  3  Online           54     30          0       1024         3          25
```

Sample Output

show chassis fpc (SRX5600 and SRX5800 devices)

```
user@host> show chassis fpc
                     Temp   CPU Utilization (%)       Memory    Utilization (%)
Slot State           (C)   Total  Interrupt         DRAM (MB) Heap     Buffer
  0  Empty
  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
```

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 (SRX1400, SRX3400, and SRX3600 devices)

```
user@host> show chassis fpc pic-status
Slot 0   Online       SRX3k SFB 12GE
  PIC 0  Online       8x 1GE-TX 4x 1GE-SFP
Slot 1   Online       SRX3k 2x10GE XFP
  PIC 0  Online       2x 10GE-XFP
Slot 3   Online       SRX1k3k 2x10GE NP-IOC
  PIC 0  Online       2x 10GE-SFP+
Slot 4   Online       SRX3k SPC
  PIC 0  Online       SPU Cp-Flow
```

```
Slot 5   Online       SRX1k3k 2x10GE NP-IOC
  PIC 0  Online       2x 10GE-SFP+
Slot 6   Online       SRX3k NPC
  PIC 0  Online       NPC PIC
Slot 7   Online       SRX1k3k 2x10GE NP-IOC
  PIC 0  Online       2x 10GE-SFP+- services-offload low-latency
```

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
```

Sample Output

show chassis fpc pic-status for HA (SRX3400 and SRX3600 devices)

```
user@host> show chassis fpc pic-status
node0:
--------------------------------------------------------------------------
Slot 0   Online       SRX3k SFB 12GE
  PIC 0  Online       8x 1GE-TX 4x 1GE-SFP
Slot 1   Online       SRX3k 2x10GE XFP
  PIC 0  Online       2x 10GE-XFP
Slot 2   Online       SRX3k 16xGE SFP
  PIC 0  Online       16x 1GE-SFP
Slot 7   Online       SRX3k SPC
  PIC 0  Online       SPU Cp-Flow
Slot 11  Online       SRX3k NPC
  PIC 0  Online       NPC PIC
Slot 12  Online       SRX3k NPC
  PIC 0  Online       NPC PIC

node1:
--------------------------------------------------------------------------
Slot 0   Online       SRX3k SFB 12GE
  PIC 0  Online      8x 1GE-TX 4x 1GE-SFP
Slot 1   Online       SRX3k 2x10GE XFP
  PIC 0  Online      2x 10GE-XFP
Slot 2   Online       SRX3k 16xGE SFP
  PIC 0  Online      16x 1GE-SFP
Slot 7   Online       SRX3k SPC
  PIC 0  Online       SPU Cp-Flow
Slot 11  Online       SRX3k NPC
  PIC 0  Online      NPC PIC
Slot 12  Online       SRX3k NPC
  PIC 0  Online      NPC PIC
```

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
```

Sample Output

show chassis fpc pic-status (SRX5600 and SRX5800 devices with NG-SPC)

```
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
```

Sample Output

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 hardware (View)

Supported Platforms J Series, LN Series, SRX Series

Syntax
show chassis hardware
<**clei-models | detail | extensive |** models | node (***node-id*** | all | local | primary)>

Release Information Command modified in Release 9.2 of Junos OS; **node** options added in Release 9.0 of Junos OS.

Description Display chassis hardware information.

Options
- **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.
  - ***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
- *Flow-Based Processing Feature Guide for Security Devices*
- *Ethernet Interfaces Feature Guide for Security Devices*
- *Junos OS Interfaces Library for Security Devices*

List of Sample Output

Output Fields Table 12 on page 173 lists the output fields for the **show chassis hardware** command. Output fields are listed in the approximate order in which they appear.

Table 12: 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.
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).
Description	Brief description of the hardware item: • Type of power supply. • Switch Control Board (SCB) • Type of Flexible PIC Concentrator (FPC), IOC, Physical Interface Card (PIC), Modular Interface Cards (MICs), and PIMs. • SRX Clustering Module • Fan Tray • For hosts, the Routing Engine type.

Sample Output

show chassis hardware

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                AB1609AA0082      SRX 3600
Midplane         REV 07   710-020310   VP8136            SRX 3600 Midplane
PEM 0            rev 05   740-027644   G087E6003S05P     AC Power Supply
PEM 1            rev 05   740-027644   G087E600AT05P     AC Power Supply
CB 0             REV 11   750-021914   AAAC9887          SRX3k RE-12-10
  Routing Engine          BUILTIN      BUILTIN           Routing Engine
  CPP                     BUILTIN      BUILTIN           Central PFE Processor
  Mezz           REV 08   710-021035   AAAD9202          SRX HD Mezzanine Card
FPC 0            REV 11   750-021882   AAAD9785          SRX3k SFB 12GE
  PIC 0                   BUILTIN      BUILTIN           8x 1GE-TX 4x 1GE-SFP
    Xcvr 8       REV 01   740-011613   PDG0UMW           SFP-SX
    Xcvr 9       REV 02   740-011613   PGJ5GJF           SFP-SX
```

```
    Xcvr 11      REV 01   740-014132   62081010          SFP-T
FPC 1            REV 10   750-016077   AAAE9989          SRX3k SPC
  PIC 0                   BUILTIN      BUILTIN           SPU Cp-Flow
FPC 2            REV 11   750-016077   AAAT8490          SRX3k SPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
FPC 5            REV 15   750-020321   AABB3820          SRX3k 2x10GE XFP
  PIC 0                   BUILTIN      BUILTIN           2x 10GE-XFP
    Xcvr 0       REV 03   740-014289   0ZT805000069      XFP-10G-SR
    Xcvr 1       REV 03   740-011571   C933BK00F         XFP-10G-SR
FPC 10           REV 12   750-043828   AAAD9501          SRX1k3k 2x10GE NP-IOC

  PIC 0                   BUILTIN      BUILTIN           2x 10GE-SFP+
Fan Tray 0       REV 06   750-021599   VR9734            SRX 3600 Fan Tray
```

Sample Output

show chassis hardware (SRX5600 and SRX5800 devices)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN10B7005AGA      SRX 5800
Midplane         REV 03   710-013698   TR0779            SRX 5800 Backplane
FPM Board        REV 03   710-014974   KC3406            Front Panel Display
PDM              Rev 03   740-013110   QCS1122504F       Power Distribution Module
PEM 1            Rev 03   740-013683   QCS1134703V       DC Power Entry Module
PEM 2            Rev 03   740-013683   QCS1134700E       DC Power Entry Module
Routing Engine 0 REV 06   740-015113   1000696955        RE-S-1300
CB 0             REV 07   710-013385   JZ3257            SRX5k SCB
FPC 3            BB-P2-39 710-020305   JS4847            SRX5k SPC
  CPU            REV 06   710-013713   KC1180            DPC PMB
  PIC 0                   BUILTIN      BUILTIN           SPU Cp
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
FPC 6            REV 03   750-020751   JT0109            SRX5k DPC 4x 10GE
  CPU            REV 06   710-013713   KC3543            DPC PMB
  PIC 0                   BUILTIN      BUILTIN           1x 10GE(LAN/WAN) RichQ
    Xcvr 0                NON-JNPR     A7C00SY           XFP-10G-SR
  PIC 1                   BUILTIN      BUILTIN           1x 10GE(LAN/WAN) RichQ
    Xcvr 0       REV 01   740-011571   C728XJ01W         XFP-10G-SR
  PIC 2                   BUILTIN      BUILTIN           1x 10GE(LAN/WAN) RichQ
  PIC 3                   BUILTIN      BUILTIN           1x 10GE(LAN/WAN) RichQ
Fan Tray 0       REV 04   740-014971   TP1432            Fan Tray
Fan Tray 1       REV 04   740-014971   TP1829            Fan Tray
```

Sample Output

show chassis hardware detail

```
user@host> show chassis hardware detail
Hardware inventory:
Item             Version  Part number  Serial number     Description
    Chassis                                JN000968AB        J4300
    Midplane         REV 05   710-010001   ad04420077
    System IO        REV 07   710-010003   AE04420393        System IO board
    Routing Engine   REV 08   750-010005   btrd43500196      RE-J.2
      ad0     488 MB  512MB CHH            504754C53A711400  Compact Flash
      ad2     488 MB  512MB CHH            504754C43A711400  Removable Compact
Flash
    FPC 2            REV 08   750-013493   NB9161            FPC
      PIC 0                                                  Integrated Services
```

```
Module
     ANNEX          REV 08   750-013493   NB9161           Integrated Services
```

Sample Output

show chassis hardware detail node 1

```
user@host> show chassis hardware detail node 1
node1:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN108C688ADB      J6350
Midplane         REV 03   710-014593   NM7516
System IO        REV 01   710-016210   NN9781            JX350 System IO
Crypto Module                                           Crypto Acceleration
Routing Engine   REV 08   710-015273   NM6569            RE-J6350-3400
  ad0     991 MB  1GB CKS              2006000000000000800 Compact Flash
FPC 0                                                   FPC
FPC 3            REV 06   750-013492   NM1294            FPC
FPC 6            REV 11   750-015153   NP8750            FPC
```

Sample Output

show chassis hardware extensive

```
user@host> show chassis hardware extensive
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN107a494ADA      J4350
 Jedec Code:   0x0000            EEPROM Version:    0x00
 P/N:          ..........        S/N:               JN107a494ADA
 Assembly ID:  0x0513            Assembly Version:  00.00
 Date:         00-00-0000        Assembly Flags:    0x00
 Version:      ..........
 ID: j4350
...
FPC 5            REV 08   750-013493   NB9161            FPC
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          750-013493        S/N:               S/N NB9161
 Assembly ID:  0x073c            Assembly Version:  01.08
 Date:         03-03-2006        Assembly Flags:    0x00
 Version:      REV 08
 ID: FPC                         FRU Model Number:  SSG-EPIM-1TX
 Board Information Record:
  Address 0x00: 34 01 05 05 02 ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 07 3c 01 08 52 45 56 20 30 38 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 31 33 34 39 33 00 00
  Address 0x20: 53 2f 4e 20 4e 42 39 31 36 31 00 00 00 03 03 07
  Address 0x30: d6 ff ff ff 34 01 05 05 02 ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 53
  Address 0x50: 53 47 2d 45 50 49 4d 2d 31 54 58 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  PIC 0                                                 Integrated Services
Module
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 Assembly ID:  0x063c            Assembly Version:  01.08
 Date:         03-03-2006        Assembly Flags:    0x00
 ID: Integrated Services Module
```

```
Board Information Record:
 Address 0x00: 34 01 05 05 02 ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
 Address 0x00: 7f b0 01 ff 06 3c 01 08 00 00 00 00 00 00 00 00
 Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x20: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 03 03 07
 Address 0x30: d6 ff ff ff 34 01 05 05 02 ff ff ff ff ff ff ff
 Address 0x40: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 53
 Address 0x50: 53 47 2d 45 50 49 4d 2d 31 54 58 00 00 00 00 00
 Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
 Address 0x70: ff ff ff ff ed c0 0e 0b 00 00 00 00 00 00 00 00
 ANNEX          REV 08   750-013493   NB9161            Integrated Services
Jedec Code:   0x7fb0            EEPROM Version:    0x01
P/N:          750-013493        S/N:               S/N NB9161
Assembly ID:  0x0808            Assembly Version:  01.08
Date:         03-03-2006        Assembly Flags:    0x00
Version:      REV 08
ID: Integrated Services         FRU Model Number:  SSG-EPIM-1TX
Board Information Record:
 Address 0x00: 34 01 05 05 02 ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
 Address 0x00: 7f b0 01 ff 08 08 01 08 52 45 56 20 30 38 00 00
 Address 0x10: 00 00 00 00 37 35 30 2d 30 31 33 34 39 33 00 00
 Address 0x20: 53 2f 4e 20 4e 42 39 31 36 31 00 00 00 03 03 07
 Address 0x30: d6 ff ff ff 34 01 05 05 02 ff ff ff ff ff ff ff
 Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 53
 Address 0x50: 53 47 2d 45 50 49 4d 2d 31 54 58 00 00 00 00 00
 Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
 Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
```

Sample Output

show chassis hardware models (SRX1400 and SRX3000 devices)

```
user@host> show chassis hardware models
Hardware inventory:
Item             Version  Part number  Serial number     FRU model number
Midplane         REV 07   710-020310   VP8136            SRX3600-CHAS
PEM 0            rev 05   740-027644   G087E6003S05P     AC Power Supply
PEM 1            rev 05   740-027644   G087E600AT05P     AC Power Supply
CB 0             REV 11   750-021914   AAAC9887          SRX3K-RE-12-10
  Routing Engine          BUILTIN      BUILTIN
  CPP                     BUILTIN      BUILTIN
FPC 0            REV 11   750-021882   AAAD9785          SRX3K-SFB-12GE
FPC 1            REV 10   750-016077   AAAE9989          SRX3K-SPC-1-10-40
FPC 2            REV 11   750-016077   AAAT8490          SRX3K-SPC-1-10-40
FPC 5            REV 15   750-020321   AABB3820          SRX3K-2XGE-XFP
FPC 10           REV 12   750-043828   AAAD9501          SRX1K3K-NP-2XGE-SFPP
Fan Tray 0       REV 06   750-021599   VR9734            SRX3600-FAN
```

show chassis hardware models (SRX5600 and SRX5800 devices)

```
user@host> show chassis hardware models
Hardware inventory:
Item             Version  Part number  Serial number     FRU model number
Midplane         REV 03   710-013698   TR0779            CHAS-BP-MX960-S
FPM Board        REV 03   710-014974   KC3406            CRAFT-MX960-S
Routing Engine 0 REV 06   740-015113   1000696955        RE-S-1300-2048-S
CB 0             REV 07   710-013385   JZ3257            SCB-MX960-S
```

Sample Output

show chassis hardware clei-models (SRX5600 and SRX5800 devices)

```
user@host> show chassis hardware clei-models
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
FPM Board        REV 02   710-017254                     CRAFT-MX480-S
Routing Engine 0 REV 06   740-015113                     RE-S-1300-2048-S
CB 0             REV 07   710-013385                     SCB-MX960-S
Fan Tray                                                 SRX5600-FAN
```

Sample Output

show chassis hardware models (SRX5600 and SRX5800 devices for next-generation SPC)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                               JN11C3B72AGA      SRX 5800
Midplane         REV 01   710-024803   ABAB6412          SRX 5800 Backplane
FPM Board        REV 01   710-024632   YS1437            Front Panel Display
PDM              Rev 03   740-013110   QCS152250EZ       Power Distribution Modu
le
PEM 0            Rev 05   740-023516   QCS1452F08U       DC Power Entry Module
PEM 1            Rev 05   740-023516   QCS1452F084       DC Power Entry Module
Routing Engine 0 REV 09   740-023530   9009080059        SRX5k RE-13-20
CB 0             REV 05   710-024802   ZJ7533            SRX5k SCB
CB 1             REV 05   710-024802   ZJ7603            SRX5k SCB
FPC 0            REV 25   750-020235   YF6692            SRX5k DPC 40x 1GE
  CPU            REV 03   710-024633   YF5990            SRX5k DPC PMB
  PIC 0                   BUILTIN      BUILTIN           10x 1GE RichQ
    Xcvr 0       REV 02   740-011613   PJ93U0P           SFP-SX
    Xcvr 1       REV 02   740-011613   PJ93U7C           SFP-SX
  PIC 1                   BUILTIN      BUILTIN           10x 1GE RichQ
  PIC 2                   BUILTIN      BUILTIN           10x 1GE RichQ
    Xcvr 0       REV 02   740-011613   PJ93U95           SFP-SX
    Xcvr 1       REV 02   740-011613   PJ93YZW           SFP-SX
  PIC 3                   BUILTIN      BUILTIN           10x 1GE RichQ
FPC 9            REV 02   750-044175   ZY2574            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
Fan Tray 0       REV 05   740-014971   VT5896            Fan Tray
Fan Tray 1       REV 05   740-014971   VT5894            Fan Tray
```

Sample Output

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 Modu
le
PEM 0            Rev 03   740-034724   QCS17020203F      PS 4.1kW; 200-240V AC i
```

```
n
PEM 1            Rev 03   740-034724   QCS17020203C      PS 4.1kW; 200-240V AC i
n
PEM 2            Rev 04   740-034724   QCS17100200A      PS 4.1kW; 200-240V AC i
n
PEM 3            Rev 03   740-034724   QCS17080200M      PS 4.1kW; 200-240V AC i
n
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   AMB0HG3           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 ethernet-switching mac-learning-log (View)

Supported Platforms LN Series, SRX100, SRX110, SRX210, SRX220, SRX240, SRX550, SRX650

Syntax **show ethernet-switching mac-learning-log**

Release Information Command introduced in Release 9.5 of Junos OS.

Description Displays the event log of learned MAC addresses.

Required Privilege Level view

Related Documentation
- *Understanding MAC Limiting*
- *Ethernet Interfaces Feature Guide for Security Devices*

Output Fields Table 13 on page 179 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 13: 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: • 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 00:00:00:00:00:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 9 mac 00:00:00:00:00:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 10 mac 00:00:00:00:00:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 11 mac 00:00:00:00:00:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 12 mac 00:00:00:00:00:00 was deleted
```

```
Wed Mar 18 08:07:05 2009
vlan_idx 13 mac 00:00:00:00:00:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 14 mac 00:00:00:00:00:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 15 mac 00:00:00:00:00:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 16 mac 00:00:00:00:00:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 4 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 6 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 7 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 9 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 10 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 11 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 12 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 13 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 14 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 15 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 16 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 5 mac 00:00:00:00:00:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 18 mac 00:00:05:00:00:05 was learned
Wed Mar 18 08:07:05 2009
vlan_idx 5 mac 00:30:48:90:54:89 was learned
Wed Mar 18 08:07:05 2009
vlan_idx 6 mac 00:00:5e:00:01:00 was learned
Wed Mar 18 08:07:05 2009
vlan_idx 16 mac 00:00:5e:00:01:08 was learned
Wed Mar 18 08:07:05 2009
vlan_idx 7 mac 00:00:5e:00:01:09 was learned
Wed Mar 18 08:07:05 2009
vlan_idx 8 mac 00:19:e2:50:ac:00 was learned
Wed Mar 18 08:07:05 2009
vlan_idx 12 mac 00:00:5e:00:01:04 was learned
[output truncated]
```

show ethernet-switching table (View)

Supported Platforms LN Series, SRX100, SRX110, SRX210, SRX220, SRX240, SRX550, SRX650

Syntax **show ethernet-switching table (brief |detail |extensive) interface** ***interface-name***

Release Information Command introduced in Release 9.5 of Junos OS.

Description Displays the Ethernet switching table.

Options

- **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

- *Port Security Overview*
- *Understanding MAC Limiting*

Output Fields Table 14 on page 181 lists the output fields for the show ethernet-switching table command. Output fields are listed in the approximate order in which they appear.

Table 14: 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: • 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 00:00:05:00:00:03 Learn 0 ge-0/0/44.0
F2 00:19:e2:50:7d:e0 Static - Router
Linux * Flood - All-members
Linux 00:19:e2:50:7d:e0 Static - Router
Linux 00:30:48:90:54:89 Learn 0 ge-0/0/47.0
T1 * Flood - All-members
T1 00:00:05:00:00:01 Learn 0 ge-0/0/46.0
T1 00:00:5e:00:01:00 Static - Router
T1 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
T1 00:19:e2:50:7d:e0 Static - Router
T10 * Flood - All-members
T10 00:00:5e:00:01:09 Static - Router
T10 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
T10 00:19:e2:50:7d:e0 Static - Router
T111 * Flood - All-members
T111 00:19:e2:50:63:e0 Learn 0 ge-0/0/15.0
T111 00:19:e2:50:7d:e0 Static - Router
T111 00:19:e2:50:ac:00 Learn 0 ge-0/0/15.0
T2 * Flood - All-members
T2 00:00:5e:00:01:01 Static - Router
T2 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
T2 00:19:e2:50:7d:e0 Static - Router
T3 * Flood - All-members
T3 00:00:5e:00:01:02 Static - Router
T3 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
T3 00:19:e2:50:7d:e0 Static - Router
T4 * Flood - All-members
T4 00:00:5e:00:01:03 Static - Router
T4 00:19:e2:50:63:e0 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 00:00:05:00:00:03 Learn 0 ge-0/0/44.0
F2 00:19:e2:50:7d:e0 Static - Router
Linux * Flood - All-members
Linux 00:19:e2:50:7d:e0 Static - Router
Linux 00:30:48:90:54:89 Learn 0 ge-0/0/47.0
T1 * Flood - All-members
T1 00:00:05:00:00:01 Learn 0 ge-0/0/46.0
T1 00:00:5e:00:01:00 Static - Router
T1 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
T1 00:19:e2:50:7d:e0 Static - Router
T10 * Flood - All-members
T10 00:00:5e:00:01:09 Static - Router
T10 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
T10 00:19:e2:50:7d:e0 Static - Router
T111 * Flood - All-members
T111 00:19:e2:50:63:e0 Learn 0 ge-0/0/15.0
T111 00:19:e2:50:7d:e0 Static - Router
T111 00:19:e2:50:ac:00 Learn 0 ge-0/0/15.0
T2 * Flood - All-members
```

```
T2 00:00:5e:00:01:01 Static - Router
T2 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
T2 00:19:e2:50:7d:e0 Static - Router
T3 * Flood - All-members
T3 00:00:5e:00:01:02 Static - Router
T3 00:19:e2:50:63:e0 Learn 0 ge-0/0/46.0
T3 00:19:e2:50:7d:e0 Static - Router
T4 * Flood - All-members
T4 00:00:5e:00:01:03 Static - Router
T4 00:19:e2:50:63:e0 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, 00:00:05:00:00:03
Interface(s): ge-0/0/44.0
Type: Learn, Age: 0, Learned: 2:03:09
F2, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Linux, *
Interface(s): ge-0/0/47.0
Type: Flood
Linux, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Linux, 00:30:48:90:54:89
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, 00:00:05:00:00:01
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
T1, 00:00:5e:00:01:00
Interface(s): Router
Type: Static
T1, 00:19:e2:50:63:e0
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
T1, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
T10, *
Interface(s): ge-0/0/46.0
Type: Flood
T10, 00:00:5e:00:01:09
Interface(s): Router
Type: Static
T10, 00:19:e2:50:63:e0
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:08
T10, 00:19:e2:50:7d:e0
```

```
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, 00:00:05:00:00:03
Interface(s): ge-0/0/44.0
Type: Learn, Age: 0, Learned: 2:03:09
F2, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Linux, *
Interface(s): ge-0/0/47.0
Type: Flood
Linux, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Linux, 00:30:48:90:54:89
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, 00:00:05:00:00:01
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
T1, 00:00:5e:00:01:00
Interface(s): Router
Type: Static
T1, 00:19:e2:50:63:e0
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
T1, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
T10, *
Interface(s): ge-0/0/46.0
Type: Flood
T10, 00:00:5e:00:01:09
Interface(s): Router
Type: Static
T10, 00:19:e2:50:63:e0
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:08
T10, 00:19:e2:50:7d:e0
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       00:00:05:00:00:05   Learn    0 ge-0/0/1.0
```

show igmp-snooping route (View)

Supported Platforms J Series, LN Series, SRX100, SRX110, SRX210, SRX220, SRX240, SRX550, SRX650

Syntax show igmp-snooping route (brief | detail | ethernet-switching | inet | vlan)

Release Information Command introduced in Release 9.5 of Junos OS.

Description Display IGMP snooping route information.

Options

- **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 *vlan-id* |*vlan-name***—(Optional) Display route information for the specified VLAN.

Required Privilege Level view

Related Documentation

- *Ethernet Interfaces Feature Guide for Security Devices*
- *Layer 2 Bridging and Transparent Mode Feature Guide for Security Devices*

Output Fields Table 15 on page 186 lists the output fields for the **show igmp-snooping route** command. Output fields are listed in the approximate order in which they appear.

Table 15: 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         224.1.1.1, *   533
Interfaces: ge-0/0/13.0, ge-0/0/1.0
v12         224.1.1.3, *   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        224.1.1.1, *       1266
Interfaces: ge-0/0/0.0
v1        224.1.1.3, *       1266
Interfaces: ge-0/0/0.0
v1        224.1.1.5, *       1266
Interfaces: ge-0/0/0.0
v1        224.1.1.7, *       1266
Interfaces: ge-0/0/0.0
v1        224.1.1.9, *       1266
Interfaces: ge-0/0/0.0
v1        224.1.1.11, *      1266
Interfaces: ge-0/0/0.0
```

show lacp interfaces

Supported Platforms LN Series, M Series, MX Series, PTX Series, T Series

Syntax
show lacp interfaces
<*interface-name*>

Release Information Command introduced in Junos OS Release 7.6.

Description Display Link Aggregation Control Protocol (LACP) information about the specified aggregated Ethernet, Fast Ethernet, or Gigabit Ethernet interface.

Options **none**—Display LACP information for all interfaces.

interface-name—(Optional) Display LACP information for the specified interface:

- Aggregated Ethernet—**ae*number***
- Fast Ethernet—**fe-*fpc/pic/port***
- Gigabit Ethernet—**ge-*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
- *Ethernet Interfaces Feature Guide for Security Devices*

List of Sample Output
show lacp interfaces (Aggregated Ethernet) on page 191
show lacp interfaces (Gigabit Ethernet) on page 191

Output Fields Table 16 on page 188 lists the output fields for the **show lacp interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 16: show lacp interfaces Output Fields

Field Name	Field Description
Aggregated interface	Aggregated interface value.

Table 16: show lacp interfaces Output Fields *(continued)*

Field Name	Field Description

LACP State	LACP state information for each aggregated interface:
	• **Role**—Role played by the interface. It can be one of the following:
	  ◦ **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 an LACP 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 LACP 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 LAC 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 16: show lacp interfaces Output Fields *(continued)*

Field Name	Field Description

LACP Protocol	LACP protocol information for each aggregated interface:
	• Link state (active or standby) indicated in parentheses next to the interface when link protection is configured.
	• **Receive State**—One of the following values:
	  • **Current**—The state machine receives an LACP PDU and enters the **Current** state.
	  • **Defaulted**—If no LACP PDU is received before the timer for the **Current** state expires a second time, the state machine enters the **Defaulted** state.
	  • **Expired**—If no LACP 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:
	  • **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:
	  • **Attached**—Multiplexer state machine initiates the process of attaching the port to the selected aggregator.
	  • **Collecting**—**Yes** indicates that the receive function of this link is enabled with respect to its participation in an aggregation. Received frames are passed to the aggregator for collection. **No** indicates the receive function of this link is not enabled.
	  • **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.
	  • **Distributing**—**Yes** indicates that the transmit function of this link is enabled with respect to its participation in an aggregation. Frames may be passed down from the aggregator's distribution function for transmission. **No** indicates the transmit function of this link is not enabled.
	  • **Waiting**—Multiplexer state machine is in a holding process, awaiting an outcome.

LACP Statistics	LACP statistics are returned when the **extensive** option is used and provides the following information: • **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 interfaces (Aggregated Ethernet)

```
user@host> show lacp interfaces ae0 extensive
 Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout    Activity

      ge-1/0/1       Actor    No   Yes    No   No   No   Yes     Fast      Active

      ge-1/0/1     Partner    No   Yes    No   No   No   Yes     Fast      Passive

      ge-1/0/2       Actor    No   Yes    No   No   No   Yes     Fast      Active

      ge-1/0/2     Partner    No   Yes    No   No   No   Yes     Fast      Passive

    LACP protocol:        Receive State    Transmit State      Mux State
      ge-1/0/1                CURRENT       Fast periodic      Collecting
distributing
      ge-1/0/2                CURRENT       Fast periodic      Collecting
distributing
      ge-1/0/1 (active)       CURRENT       Fast periodic      Collecting
distributing
      ge-1/0/2 (standby)      CURRENT       Fast periodic      WAITING
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      ge-1/0/1                   0           0            0   0
      ge-1/0/2                   0           0            0   0
```

show lacp interfaces (Gigabit Ethernet)

```
user@host> show lacp interfaces ge-0/3/0
Aggregated interface: ae0
    LACP State:       Role   Exp   Def  Dist  Col  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 interfaces (View)

Supported Platforms J Series, LN Series, SRX Series

Syntax **show lacp interfaces** ***interface-name***

Release Information Command modified in Release 10.2 of Junos OS.

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 ***interface-name***—(Optional) Display LACP information for the specified interface:

- Aggregated Ethernet—**ae*number***
- Redundant Ethernet—**reth*number***
- Gigabit Ethernet—**ge-*fpc/pic/port***
- 10-Gigabit Ethernet—**xe*number***

Required Privilege Level view

Related Documentation

- *Ethernet Interfaces Feature Guide for Security Devices*

List of Sample Output

show lacp interfaces (Aggregated Ethernet) on page 194
show lacp interfaces (Redundant Ethernet) on page 195
show lacp interfaces (Gigabit Ethernet) on page 195

Output Fields Table 17 on page 192 lists the output fields for the **show lacp interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 17: show lacp interfaces Output Fields

Field Name	Field Description
Aggregated interface	Aggregated interface value.

Table 17: show lacp interfaces Output Fields *(continued)*

Field Name	Field Description

LACP State	LACP state information for each aggregated interface:
	• **Role**—Role played by the interface. It can be one of the following:
	  ◦ **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 17: show lacp interfaces Output Fields *(continued)*

Field Name	Field Description

LACP Protocol	LACP protocol information for each aggregated interface:
	• Link state (active or standby) indicated in parentheses next to the interface when link protection is configured.
	• **Receive State**—One of the following values:
	  • **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:
	  • **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:
	  • **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  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}
```

show lacp interfaces (Gigabit Ethernet)

```
user@host> show lacp interfaces ge-0/3/0
Aggregated interface: ae0
    LACP State:       Role   Exp   Def  Dist  Col  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 J Series, LN Series, QFX Series, SRX Series

Syntax **show lacp statistics interfaces** *interface-name*

Release Information Command modified in Release 10.2 of Junos OS.
Command introduced in Release 11.1 of Junos OS for the QFX Series.

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
- *Ethernet Port Switching Feature Guide for Security Devices*

List of Sample Output
show lacp statistics interfaces on page 196
show lacp statistics interfaces (QFX Series) on page 197
show lacp statistics interfaces (QFabric Switches) on page 197

Output Fields Table 18 on page 196 lists the output fields for the **show lacp statistics interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 18: show lacp statistics interfaces Output Fields

Field Name	Field Description
Aggregated interface	Aggregated interface value.
LACP Statistics	LACP statistics provide the following information: • **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 lacp statistics interfaces (QFX Series)

```
user@host> show lacp statistics interfaces ae0
Aggregated interface: ae0
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      xe-0/0/2                1352        2035            0            0
      xe-0/0/3                1352        2056            0            0
```

show lacp statistics interfaces (QFabric Switches)

```
user@host> show lacp statistics interfaces nodegroup1:ae0
Aggregated interface: nodegroup1:ae0
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      node1:xe-0/0/2              1352        2035            0            0
      node2:xe-0/0/3              1352        2056            0            0
```

show oam ethernet link-fault-management

Supported Platforms LN Series, SRX100, SRX210, SRX220, SRX240, SRX550, SRX650

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
- *Ethernet Port Switching Feature Guide for Security Devices*
- Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways on page 31
- Example: Configuring Ethernet OAM Link Fault Management on page 95

List of Sample Output
show oam ethernet link-fault-management brief on page 202
show oam ethernet link-fault-management detail on page 202

Output Fields Table 19 on page 198 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 19: show oam ethernet link-fault-management Output Fields

Field Name	Field Description	Level of Output
Status	Status of the established link. • **Fail**—A link fault condition exists. • **Running**—A link fault condition does not exist.	All levels
Discovery state	State of the discovery mechanism: • **Passive Wait** • **Send Any** • **Send Local Remote** • **Send Local Remote Ok**	All levels
Peer address	Address of the OAM peer.	All levels

Table 19: show oam ethernet link-fault-management Output Fields *(continued)*

Field Name	Field Description	Level of Output
Flags	Information about the interface. • **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	Remote entity information. • **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**
Loopback control	Number of loopback control PDUs received.	**detail**

Table 19: show oam ethernet link-fault-management Output Fields *(continued)*

Field Name	Field Description	Level of Output
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 19: 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: 00:90:69:72:2c:83
  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: 00:90:69:0a:07:14
    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 LN Series, SRX210, SRX220, SRX240, SRX550, SRX650

Syntax show poe controller

Release Information Command introduced in Release 9.5 of Junos OS.

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
- *Ethernet Interfaces Feature Guide for Security Devices*
- *Junos OS Interfaces Library for Security Devices*

Output Fields Table 20 on page 203 lists the output fields for the **show poe controller** command. Output fields are listed in the approximate order in which they appear.

Table 20: 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
```

PART 4

Index

Index

Symbols

A

B

C

D

E

T

V