Chapter 5

Configuring Ethernet Interfaces

This chapter describes how to configure Ethernet interfaces on E-series routers.

Most of the procedures described here do not apply to the Fast Ethernet management port on the SRP module. You can, however, select and display statistics for that port by using commands described in this chapter. For information about managing the Fast Ethernet port on the SRP module, see *JUNOSe System Basics Configuration Guide, Chapter 6, Managing Modules*.

This chapter contains the following sections:

Ethernet Overview

Ethernet modules support the following Ethernet interfaces:

- Fast Ethernet
- Gigabit Ethernet
- 10-Gigabit Ethernet
- IEEE 802.3ad link aggregation group (LAG)

This section describes features that are available with Ethernet interfaces.

Ethernet modules use the Address Resolution Protocol (ARP) to obtain MAC addresses for outgoing Ethernet frames. For more information about ARP, see *JUNOSe IP, IPv6, and IGP Configuration Guide, Chapter 1, Configuring IP*.

NOTE: Read *Configuration Tasks for Ethernet* on page 183 before you begin to configure an Ethernet interface.

Features

Ethernet interfaces support the following features:

- Routing of IP packets.
- Quality of Service (QoS) classification.
- High-density Ethernet. (For information, see *High-Density Ethernet* on page 181.*)*
- Virtual LAN (VLAN) configurations. (For information, see *Configuring VLANs* on page 187.)
- Stacked Virtual LAN (S-VLAN) configurations. (For information, see *Configuring S-VLANs* on page 197.)
- Configurations with higher-level protocols. (For information, see *Configuring Higher-Level Protocols over Ethernet* on page 225.)
- Layer 2 Tunneling Protocol (L2TP. For information, see *L2TP* on page 165.)
- Multinetting. (For information, see *Multinetting* on page 165.)

L2TP

Most Ethernet interfaces support L2TP. To use L2TP, you must first create a PPP interface. See *JUNOSe Broadband Access Configuration Guide, Chapter 10, L2TP Overview* for information about configuring L2TP.

Multinetting

Ethernet interfaces, except for bridged Ethernet interfaces, support multinetting; that is, adding more than one IP address to an IP interface. If you want to add multiple IP addresses to a single IP interface, use the **ip address** command with the **secondary** keyword, which is described in *JUNOSe IP, IPv6, and IGP Configuration Guide, Chapter 1, Configuring IP*.

Ethernet Interface Platform Considerations

You can configure Ethernet interfaces on the following E-series routers:

- E120 router
- E320 router
- ERX-1440 router
- ERX-1410 router
- ERX-710 router
- ERX-705 router
- ERX-310 router

This section describes the line modules and I/O modules that support Ethernet interfaces.

For detailed information about the modules that support Fast Ethernet and Gigabit Ethernet interfaces on the ERX-14xx models, ERX-7xx models, and the ERX-310 router:

- See *ERX Module Guide, Table 1, Module Combinations* for detailed specifications of these modules.
- See *ERX Module Guide, Appendix A, Module Protocol Support* for information about the protocols and applications that Ethernet modules support.

For detailed information about the modules that support Gigabit Ethernet and 10-Gigabit Ethernet interfaces on the E120 router and the E320 router:

- See *E120 and E320 Module Guide, Table 1, Modules and IOAs* for detailed specifications of these modules.
- See *E120 and E320 Module Guide, Appendix A, IOA Protocol Support* for information about the protocols and applications that Ethernet modules support.

Numbering Scheme

When configuring or managing an interface, you must know the numbering scheme for identifying an interface. The numbering scheme depends on the type of E-series router that you have.

ERX-7xx Models, ERX-14xx Models, and the ERX-310 Router

Use the *slot/port* [.*subinterface*] format to identify Ethernet interfaces and subinterfaces.

- *slot*—Number of the slot in which the line module resides in the chassis.

  In ERX-7xx models, line module slots are numbered 2-6; slots 0 and 1 are reserved for SRP modules. In ERX-14xx models, line module slots are numbered 0–5 and 8–13; slots 6 and 7 are reserved for SRP modules. In an ERX-310 router, line module slots are numbered 1–2; slot 0 is reserved for the SRP module.

- *port*—Number of the port on the I/O module.

  On the OC3-2 GE APS I/O module, you can configure only a Gigabit Ethernet interface in port 2; ports 0 and 1 are reserved for OC3/STM1 ATM interfaces.

- *subinterface*—Subinterface number of the protocol or VLAN subinterface.

For information about installing line modules and I/O modules in ERX routers, see *ERX Hardware Guide, Chapter 4, Installing Modules*.

E120 Router and E320 Router

Use the *slot/adapter/port* [.*subinterface*] format to identify Ethernet interfaces and subinterfaces.

- *slot*—Number of the slot in which the line module resides in the chassis.

  In the E120 router, line module slots are numbered 0–5. In the E320 router, line module slots are numbered 0–5 and 11–16. For both routers, slots 6 and 7 are reserved for SRP modules; slots 8–10 are reserved for switch fabric modules (SFMs).

- *adapter*—Number of the bay in which the I/O adapter (IOA) resides.

  This identifier applies to the E120 and E320 routers only. In the software, adapter 0 identifies the right IOA bay (E120 router) and the upper IOA bay (E320 router); adapter 1 identifies the left IOA bay (E120 router) and the lower IOA bay (E320 router).

- *port*—Number of the port on the IOA

- *subinterface*—Subinterface number of the protocol or VLAN subinterface

For information about installing line modules and IOAs in the E120 and E320 routers, see *E120 and E320 Hardware Guide, Chapter 4, Installing Modules*.

Interface Specifier

The configuration task examples in this chapter use the format for ERX-7xx models, ERX-14xx models, and the ERX-310 router to specify an Ethernet interface. (The format is described in *Numbering Scheme* on page 166.)

For example, the following command specifies a Fast Ethernet interface on port 0 of an I/O module in slot 4.

host1(config)#**interface fastEthernet 4/0**

When you configure a Gigabit Ethernet interface or a 10-Gigabit Ethernet interface on E120 or E320 routers, you must include the adapter identifier as part of the interface specifier. For example, the following command specifies a Gigabit Ethernet interface on port 0 of the IOA installed in the upper adapter bay of slot 3.

host1(config)#**interface gigabitEthernet 3/0/0**

For more information about interface types and specifiers on E-series models, see *Interface Types and Specifiers* in *JUNOSe Command Reference Guide, About This Guide*.

Fast Ethernet I/O Modules

Fast Ethernet interfaces are supported on two I/O modules: the FE-8 I/O module and the FE-8 SFP I/O module.

FE-8 I/O Module

ERX-7xx models, ERX-14xx models, and the ERX-310 router all support the FE-8 I/O module.

An FE-8 I/O module accepts up to eight RJ-45 connectors.

FE-8 SFP I/O Module

ERX-7xx models, ERX-14xx models, and the ERX-310 router all support the FE-8 SFP I/O module.

The FE-8 SFP I/O module uses a range of small form-factor pluggable transceivers (SFPs) to support different optical modes and cabling distances. The I/O module supports up to eight LC-style fiber-optic connectors.

Unlike all other Fast Ethernet and Gigabit Ethernet I/O modules, the FE-8 SFP I/O module does not support automatic negotiation of the line speed and duplex mode by the router. For more information, see **duplex** on page 184 and **speed** on page 187.

Gigabit Ethernet I/O Modules and IOAs

Gigabit Ethernet interfaces are supported on the following modules:

- GE I/O module
- GE-2 SFP I/O module
- GE-8 I/O module
- OC3-2 GE APS I/O module
- ES2-S1 GE-4 IOA
- ES2-S1 GE-8 IOA
- ES2-S3 GE-20 IOA

GE I/O Module

ERX-7xx models, ERX-14xx models, and the ERX-310 router all support the GE I/O module.

You can pair any of the following types of GE I/O modules with the GE/FE line module:

- The GE I/O SFP module uses a range of SFPs to support different optical modes and cabling distances. The I/O module accepts up to two pairs (Tx/Rx) of LC-style fiber-optic connectors that support different optical modes and cabling distances. Alternatively, the I/O module supports up to two pairs of copper SFPs.
- The GE I/O MULTI MODE module accepts up to two pairs (Tx/Rx) of SC-style fiber-optic connectors.
- The GE I/O SINGLE MODE module accepts up to two pairs (Tx/Rx) of SC-style fiber-optic connectors.

The GE I/O module has two ports: one port (port 0) is active (also known as *primary*), and the other port (port 0R) is redundant. If the active port fails, the redundant port automatically becomes active.

You can configure only port 0 for a Gigabit Ethernet interface; you cannot configure redundant port 0R. Cabling both ports provides a redundant path to the Gigabit Ethernet interface.

GE-2 SFP I/O Module

The ERX-1440 router and the ERX-310 router both support the GE-2 SFP I/O module. Other E-series routers do not support the GE-2 SFP I/O module. The GE-2 SFP I/O module was previously called the 2XGE APS I/O module.

The GE-2 SFP I/O module pairs with the GE-2 line module and the GE-HDE line module. You can install the GE-2 line module or the GE-HDE line module and its corresponding GE-2 SFP I/O module in slot 1 or slot 2 of an ERX-310 router or in any slot of an ERX-1440 router.

The GE-2 SFP I/O module can use either fiber-optic or copper SFPs. The I/O module accepts up to two pairs (Tx/Rx) of LC-style fiber-optic connectors that support different optical modes and cabling distances. Alternatively, the I/O module supports up to two pairs of copper SFPs.

Ports on GE-2 SFP I/O Module

The GE-2 SFP I/O module has four ports: two ports (port 0 and port 1) are active (also known as *primary*), and the other two ports (port 0R and port 1R) are redundant. If an active port fails, its corresponding redundant port automatically becomes active.

You can configure only port 0 and port 1 for a Gigabit Ethernet interface; you cannot configure redundant ports 0R and 1R. Cabling an active port and its corresponding redundant port (that is, port 0 and port 0R, or port 1 and port 1R) provides a redundant path to the Gigabit Ethernet interface.

Bandwidth and Line Rate Considerations

When the GE-2 line module or the GE-HDE line module is installed in the ERX-1440 router, it delivers full bandwidth of 2 GB per port only when installed in slot 2 or slot 4, and when the SRP-40G+ module is used in the router. When the module is installed in any other ERX-1440 slot, it delivers a maximum bandwidth of 2 GB per line module (1 GB maximum at the ingress and 1 GB maximum at the egress). Therefore, of the maximum 24 possible ports for the module in an ERX-1440 chassis (that is, two ports in each of 12 slots), full bandwidth is delivered only on a maximum of four ports (those in slots 2 and 4).

When the GE-2 line module or the GE-HDE line module is installed in either the ERX-1440 router or the ERX-310 router and both ports are active, line rate performance is achieved only with packets that are 174 bytes or larger. The module might not achieve line rate with packets that are smaller than 174 bytes.

GE-8 I/O Module

The ERX-1440 router and the ERX-310 router both support the GE-8 I/O module. Other E-series routers do not support the GE-8 I/O module.

The GE-8 I/O module pairs with the GE-HDE line module to provide Gigabit Ethernet operation through eight line interfaces.

NOTE: The GE-8 I/O module has a logical port, numbered port 8, that is reserved for the hardware multicast packet replication feature. For more information, see *JUNOSe Multicast Routing Configuration Guide, Chapter 1, Configuring IPv4 Multicast* and *JUNOSe Multicast Routing Configuration Guide, Chapter 5, Configuring IPv6 Multicast*.

You can install the GE-HDE line module and its corresponding GE-8 I/O module in slot 1 or slot 2 of an ERX-310 router or in any slot of an ERX-1440 router.

The GE-8 I/O module can use either fiber-optic or copper SFPs. The I/O module accepts up to eight pairs (Tx/Rx) of LC-style fiber-optic connectors that support different optical modes and cabling distances. Alternatively, the I/O module supports up to eight pairs of copper SFPs.

Bandwidth and Line Rate Considerations

When the GE-HDE line module is installed in the ERX-1440 router, it delivers full bandwidth of 4 GB per line module (2 GB at the ingress and 2 GB at the egress) only when installed in slot 2 or slot 4, and when the SRP-40G+ module is used in the router. When the module is installed in any other ERX-1440 slot, it delivers a maximum bandwidth of 2 GB per line module (1 GB maximum at the ingress and 1 GB maximum at the egress). Therefore, of the maximum 96 possible ports for the module in an ERX-1440 chassis (that is, eight ports in each of 12 slots), full bandwidth is delivered only on a maximum of 16 ports (those in slots 2 and 4).

When the GE-HDE line module is installed in either the ERX-1440 router or the ERX-310 router and all ports are active, line rate performance is achieved only with packets that are 174 bytes or larger. The module might not achieve line rate with packets that are smaller than 174 bytes.

Table 15 lists the average data rate on the GE-HDE line module and GE-8 I/O module combination when installed in an ERX-310 router or in slots 2 or 4 of an ERX-1440 router.

Table 15: Average Data Rate for ERX-310 Router or in Slots 2 or 4 of an ERX-1440 Router

Port Combination	Average Data Rate per GE-8 I/O Module (> 174 Byte Packets)	Average Data Rate per GE-HDE Line Module
Ports 1–8	250 Mbps per port	250 Mbps per port
Any four ports	500 Mbps per port	500 Mbps per port
Any two ports	1 Gbps per port	1 Gbps per port

Table 16 lists the average data rate on the GE-HDE line module and GE-8 I/O module combination when installed in all other slots on the ERX-1440.

Table 16: Average Data Rate When Installed in All Other Slots on an ERX-1440 Router

Port Combination	Average Data Rate per GE-8 I/O Module (> 174 Byte Packets)	Average Data Rate per GE-HDE Line Module
Ports 1–8	125 Mbps per port	125 Mbps per port
Any four ports	250 Mbps per port	250 Mbps per port
Any two ports	500 Mbps per port	500 Mbps per port

Managing High-Density Ethernet

The overall data rate for the GE-HDE line module is 2 Gbps; therefore, the I/O module becomes highly oversubscribed because of the wire rate of the line module. The data rate of the GE-8 I/O module is limited with larger frames, and the packet rate is limited with smaller frames.

Currently, flow control using MAC pause frames is disabled on the GE-8 I/O module. The I/O module does not transmit or receive pause frames.

For more information about high-density Ethernet, see *High-Density Ethernet* on page 181.

OC3-2 GE APS I/O Module

ERX-7xx models, ERX-14xx models, and the ERX-310 router all support the OC3-2 GE APS I/O module.

The OC3-2 GE APS I/O module pairs with the OC3/STM1 GE/FE line module to provide Gigabit Ethernet operation through one line interface and OC3 STM1 ATM operation through two line interfaces.

The OC3-2 GE APS I/O module uses a range of SFPs to support different optical modes and cabling distances, and accepts up to three LC-style fiber-optic or copper SFPs. You can configure only port 2 for Gigabit Ethernet interfaces; port 0 and port 1 are reserved for OC3/STM1 ATM interfaces.

For more information about configuring OC3/STM-1 ATM interfaces on this I/O module, see *OC3/STM1 GE/FE Line Module* on page 73.

NOTE: The OC3-2 GE APS I/O module does not support APS in the current release.

ES2-S1 GE-4 IOA

The E120 router and the E320 router support the ES2-S1 GE-4 IOA. Other E-series routers do not support the ES2-S1 GE-4 IOA.

The ES2-S1 GE-4 IOA pairs with the ES2 4G line module (LM). For more information about the ES2 4G LM, see *ES2 4G Line Module* on page 74.

The ES2-S1 GE-4 IOA is offered in a half-height size that enables you to configure it in one of two IOA bays that are available for each slot. You can install the ES2-S1 GE-4 IOA in only one of the IOA bays per slot.For more information about installing IOAs, see the *E120 and E320 Hardware Guide*.

The ES2-S1 GE-4 IOA has four ports. The IOA can use either fiber-optic or copper SFPs. The IOA accepts up to four pairs (Tx/Rx) of LC-style fiber-optic connectors that support different optical modes and cabling distances. Alternatively, the IOA supports up to eight pairs of copper SFPs.

The ES2-S1 GE-4 IOA does not support port redundancy.

ES2-S1 GE-8 IOA

The E120 router and the E320 router support the ES2-S1 GE-8 IOA. Other E-series routers do not support the ES2-S1 GE-8 IOA.

The ES2-S1 GE-8 IOA is offered in a half-height size that enables you to configure it in either of the two IOA bays that are available for each slot. You can install the ES2-S1 GE-8 IOA in both IOA bays. For more information about installing IOAs, see the *E120 and E320 Hardware Guide*.

The ES2-S1 GE-8 IOA has eight ports. The IOA can use either fiber-optic or copper SFPs. The IOA accepts up to four pairs (Tx/Rx) of LC-style fiber-optic connectors that support different optical modes and cabling distances. Alternatively, the IOA supports up to eight pairs of copper SFPs.

NOTE: The ES2-S1 GE-8 IOA has a logical port, numbered port 8, that is reserved for the hardware multicast packet replication feature. For more information, see *JUNOSe Multicast Routing Configuration Guide, Chapter 1, Configuring IPv4 Multicast* and *JUNOSe Multicast Routing Configuration Guide, Chapter 5, Configuring IPv6 Multicast*.

The ES2-S1 GE-8 IOA pairs with either the ES2 4G line module (LM) and the ES2 10G LM to provide access services.

ES2 4G LM Combination

For more information about the ES2 4G LM, see *ES2 4G Line Module* on page 74.

When paired with the ES2 4G LM, you can combine the ES2-S1 GE-8 IOA in a slot with one of the following IOA types:

- ES2-S1 OC3-8 ATM IOA
- ES2-S1 OC12-2 ATM IOA
- ES2-S1 OC12-2 POS IOA

Bandwidth and Line Rate Considerations

Table 17 lists the average data rate on the ES2-S1 GE-8 IOA when installed in E120 and E320 routers with one ES2 4G LM installed.

Table 17: Average Data Rate for One ES2-S1 GE-8 IOA Installed with an ES2 4G LM

	100 Gbps Configuration (E320 Router)		120 Gbps and 320 Gbps Configurations (E120 and E320 Routers)	
Port Combination	**Average Data Rate per GE-8 IOA (> 128 Byte Packets)**	**Average Data Rate per ES2 4G LM**	**Average Data Rate per GE-8 IOA (> 128 Byte Packets)**	**Average Data Rate per ES2 4G LM**
All eight ports	412.5 Mbps per port	412.5 Mbps per port	475 Mbps per port	475 Mbps per port
Any four ports	825 Mbps per port	825 Mbps per port	950 Mbps per port	950 Mbps per port
Any two ports	1 Gbps per port	1 Gbps per port	1 Gbps per port	1 Gbps per port

Table 18 lists the average data rate on two ES2-S1 GE-8 IOAs when installed in E120 and E320 routers with one ES2 4G LM installed.

Table 18: Average Data Rate for Two ES2-S1 GE-8 IOAs Installed with an ES2 4G LM

	100 Gbps Configuration (E320 Router)		120 Gbps and 320 Gbps Configurations (E120 and E320 Routers)	
Port Combination	**Average Data Rate per GE-8 IOA (> 128 Byte Packets)**	**Average Data Rate per ES2 4G LM**	**Average Data Rate per GE-8 IOA (> 128 Byte Packets)**	**Average Data Rate per ES2 4G LM**
All sixteen ports	206.25 Mbps per port	206.25 Mbps per port	237.5 Mbps per port	237.5 Mbps per port
Any eight ports	412.5 Mbps per port	412.5 Mbps per port	475 Mbps per port	475 Mbps per port
Any four ports	825 Mbps per port	825 Mbps per port	950 Mbps per port	950 Mbps per port
Any two ports	1 Gbps per port	1 Gbps per port	1 Gbps per port	1 Gbps per port

Table 19 lists the average data rate when combining an ES2-S1 GE-8 IOA in one adapter bay with the ES2-S1 OC3-8 ATM IOA, or the ES2-S1 OC12-2 ATM IOA, or the ES2-S1 OC12-2 POS IOA in another adapter bay. Because the OC3/STM1 and OC12/STM4 IOAs use less than half of the full bandwidth of the ES2 4G LM, the router allocates these IOAs as much bandwidth as they can use. The ES2-S1 GE-8 IOA uses any remaining bandwidth.

Each OC12/STM4 port has a maximum theoretical bandwidth of 622 Mbps. Each OC3/STM1 port has a maximum theoretical bandwidth of 155 Mbps. Therefore, the OC12/STM4 IOAs have a maximum theoretical bandwidth of 1.244 Gbps and the OC3/STM1 IOA has an maximum theoretical bandwidth of 1.244 Gbps.

Table 19: Average Data Rate for ES2-S1 GE-8 IOA Combined with Other IOA Types in Same Slot

		100 Gbps Configuration (E320 Router)	120 Gbps and 320 Gbps Configurations (E120 and E320 Routers)
Average Data Rate per GE-8 IOA (> 128 Byte Packets)	**Average Data Rate per OC12/STM4 IOA**	**Average Data Rate per ES2 4G LM**	**Average Data Rate per ES2 4G LM**
257 Mbps per port (Ports 0–7)	622 Mbps per port (Ports 0 and 1)	GE-8 IOA—257 Mbps per port OC12/STM4—622 Mbps per port	GE-8 IOA—319.5 Mbps per port OC12/STM4—622 Mbps per port
334.75 Mbps per port (Ports 0–7)	622 Mbps (Port 1)	GE-8 IOA—334.75 Mbps per port OC12/STM4 IOA—622 Mbps for port 1	GE-8 IOA—397.25 Mbps per port OC12/STM4 IOA—622 Mbps for port 1
387.5 Mbps per port (Ports 0–7)	100 Mbps per port (Ports 0 and 1)	GE-8 IOA—387.5 Mbps per port OC12/STM4 IOA—100 Mbps per port	GE-8 IOA—450 Mbps per port OC12/STM4 IOA—100 Mbps per port

Managing High-Density Ethernet

With a 100 Gbps fabric configuration, the overall data rate for the ES2 4G LM with ES2-S1 GE-8 IOAs is 3.3 Gbps. With a 120 Gbps fabric configuration or a 320 Gbps fabric configuration, the overall data rate for the ES2 4G LM with ES2-S1 GE-8 IOAs is 3.8 Gbps. In both configurations, the line module becomes highly oversubscribed because of the IOA available wire rate. When paired with the ES2 4G LM, the data rate of the ES2-S1 GE-8 IOA is bandwidth limited with larger frames, and the packet rate is limited with smaller frames.

NOTE: The overall data rate of the ES2-S1 GE-8 IOA is 0.1 Gbps less than other IOAs that pair with the ES2 4G LM because of fair bandwidth allocation across the eight ports.

Currently, flow control using MAC pause frames is disabled on the ES2-S1 GE-8 IOA. The IOA does not transmit or receive pause frames.

For more information about high-density Ethernet on E-series routers, see *High-Density Ethernet* on page 181.

ES2 10G LM Combination

When paired with the ES2 10G LM, you can only combine the ES2-S1 GE-8 IOA in a slot with another ES2-S1 GE-8 IOA.

With a 100 Gbps fabric configuration, the E320 router can accommodate up to 2 combinations of ES2 10G LMs and ES2-S1 GE-8 IOAs. You must install a combination in either of the turbo slots (slot 2 or slot 4). The 100 Gbps allocates 10 Gbps of overall bandwidth to each of these slots.

With a 120 Gbps fabric configuration, the E120 router can accommodate up to 6 combinations of ES2 10G LMs and ES2-S1 GE-8 IOAs. You can install a combination in any of the line module slots, each of which are allocated 10 Gbps of overall bandwidth.

With a 320 Gbps fabric configuration, the E320 router can accommodate up to 12 combinations of ES2 10G LMs and ES2-S1 GE-8 IOAs. You can install a combination in any of the line module slots, each of which are allocated 10 Gbps of overall bandwidth.

Bandwidth and Line Rate Considerations

Each ES2-S1 GE-8 IOA is connected to the ES2 10G LM through a 5 Gbps bus. Therefore, the aggregate bandwidth of any IOA is limited to 5 Gbps.

Table 20 lists the average data rate on the ES2-S1 GE-8 IOA when installed in E120 and E320 routers with one ES2 10G LM installed.

Table 20: Average Data Rate for One ES2-S1 GE-8 IOA Installed with an ES2 10G LM

	100 Gbps, 120 Gbps, or 320 Gbps Configuration
Port Combination	**Average Data Rate per GE-8 IOA (> 128 Byte Packets)**
Any five ports	1 Gbps per port
All eight ports	625 Mbps per port

Table 21 lists the average data rate of two ES2-S1 GE-8 IOAs when installed in E120 and E320 routers with one ES2 10G LM installed.

Table 21: Average Data Rate for Two ES2-S1 GE-8 IOAs Installed with an ES2 10G LM

	100 Gbps, 120 Gbps, or 320 Gbps Configuration
Port Combination	**Average Data Rate per GE-8 IOA (> 128 Byte Packets)**
All sixteen ports	625 Mbps per port
Any five ports on each IOA	1 Gbps per port

Managing High-Density Ethernet

When installed in an E120 router or an E320 router with any SRP module combination, the overall data rate for the ES2 10G LM with one ES2-S1 GE-8 IOA is limited to 5 Gbps. The overall data rate for the ES2 10G LM with two ES2-S1 GE-8 IOAs is limited to 10 Gbps. In all configurations, the line module can become oversubscribed because of the IOA available wire rate (8 Gbps).

Currently, flow control using MAC pause frames is disabled on the ES2-S1 GE-8 IOA. The IOA does not transmit or receive pause frames.

For more information about high-density Ethernet on E-series routers, see *High-Density Ethernet* on page 181.

ES2-S3 GE-20 IOA

The E120 router and the E320 router support the ES2-S3 GE-20 IOA. Other E-series routers do not support the ES2-S3 GE-20 IOA.

The ES2-S3 GE-20 IOA pairs with the ES2 10G LM to provide Gigabit Ethernet operation through 20 line interfaces.

The ES2-S3 GE-20 IOA is offered in a full-height size that uses both adapter bays. The IOA is identified by the software as adapter bay 0. For more information about installing IOAs, see the *E120 and E320 Hardware Guide*.

The IOA can use either fiber-optic or copper SFPs. The IOA accepts up to four pairs (Tx/Rx) of LC-style fiber-optic connectors that support different optical modes and cabling distances.

The ES2-S3 GE-20 IOA does not support port redundancy.

ES2 10G LM Combination

With a 100 Gbps fabric configuration, the E320 router can accommodate up to 2 combinations of ES2 10G LMs and ES2-S3 GE-20 IOAs. You must install a combination in either of the turbo slots (slot 2 or slot 4). The 100 Gbps allocates 10 Gbps of overall bandwidth to each of these slots.

With a 120 Gbps fabric configuration, the E120 router can accommodate up to 6 combinations of ES2 10G LMs and ES2-S3 GE-20 IOAs. You can install a combination in any of the line module slots, each of which are allocated 10 Gbps of overall bandwidth.

With a 320 Gbps fabric configuration, the E320 router can accommodate up to 12 combinations of ES2 10G LMs and ES2-S3 GE-20 IOAs. You can install a combination in any of the line module slots, each of which are allocated 10 Gbps of overall bandwidth.

Bandwidth and Line Rate Considerations

Table 20 lists the average data rate on the ES2-S3 GE-20 IOA when installed in E120 and E320 routers with one ES2 10G LM installed.

Table 22: Average Data Rate for One ES2-S3 GE-20 IOA Installed with an ES2 10G LM

	100 Gbps, 120 Gbps, or 320 Gbps Configuration
Port Combination	Average Data Rate per GE-20 IOA (> 128 Byte Packets)
Any 10 ports	1 Gbps per port
All 20 ports	500 Mbps per port

Managing High-Density Ethernet

When installed in an E120 router or an E320 router with any SRP module combination, the overall data rate for the ES2 10G LM with one ES2-S3 GE-20 IOA is limited to 10 Gbps. The line module can become oversubscribed because of the IOA available wire rate (20 Gbps).

Currently, flow control using MAC pause frames is disabled on the ES2-S3 GE-20 IOA. The IOA does not transmit or receive pause frames.

For more information about high-density Ethernet on E-series routers, see *High-Density Ethernet* on page 181.

10-Gigabit Ethernet IOAs

10-Gigabit Ethernet interfaces are supported on the ES2-S1 10GE IOA and the ES2-S1 10GE PR IOA. For more information about 10-Gigabit Ethernet, see IEEE Standard 802.3ae.

ES2-S1 10GE IOA

The E120 router and the E320 router support the ES2-S1 10GE IOA. Other E-series routers do not support the ES2-S1 10GE IOA.

The ES2-S1 10GE IOA pairs with the ES2 4G LM to provide a 10-Gigabit Ethernet interface. For more information about the ES2 4G LM, see *ES2 4G Line Module* on page 74.

The ES2-S1 10GE IOA is offered in a full-height size that uses both adapter bays. The IOA is identified by the software as adapter bay 0. For more information about installing IOAs, see the *E120 and E320 Hardware Guide*.

The ES2-S1 10GE IOA has one port, uses a range of 10-gigabit small form-factor pluggable transceivers (XFPs) to support different optical modes and cabling distances, and accepts one LC-style fiber-optic connector.

Managing High-Density Ethernet

With a 100 Gbps fabric configuration, the overall data rate for the ES2 4G LM with the ES2-S1 10GE IOA is 3.4 Gbps for large packets. With a 120 Gbps or a 320 Gbps fabric configuration, the overall data rate for the ES2 4G LM with the ES2-S1 10GE IOA is 3.9 Gbps for large packets. In all configurations, the line module becomes highly oversubscribed because of the available wire rate on the IOA. When paired with the ES2 4G LM, the data rate of the ES2-S1 10GE IOA is bandwidth limited with larger frames, and the packet rate is limited with smaller frames.

Currently, flow control using MAC pause frames is disabled on the ES2-S1 10GE IOA. The IOA does not transmit or receive pause frames.

For more information about high-density Ethernet on E-series routers, see *High-Density Ethernet* on page 181.

ES2-S2 10GE PR IOA

The E120 router and the E320 router support the ES2-S2 10GE PR IOA. Other E-series routers do not support the ES2-S2 10GE PR IOA.

The ES2-S2 10GE PR IOA is offered in a full-height size that uses both adapter bays. The IOA is identified by the software as adapter bay 0. For more information about installing IOAs, see the *E120 and E320 Hardware Guide*.

The ES2-S2 10GE PR IOA has one port, uses a range of XFPs to support different optical modes and cabling distances, and accepts 10 LC-style fiber-optic connectors.

The single port on the ES2-S2 10GE PR IOA has a redundant port. If the active port fails, the redundant port automatically becomes active. You can configure only the active port for a 10-Gigabit Ethernet interface; you cannot configure the redundant port. Cabling both ports provides a redundant path to the 10-Gigabit Ethernet interface.

The ES2-S2 10GE PR IOA pairs with the ES2 10G Uplink LM to provide uplink services or the ES2 10G LM to provide access services.

ES2 10G Uplink LM Combination

With a 100 Gbps fabric configuration, the E320 router can accommodate up to 2 combinations of ES2 10G Uplink LMs and ES2-S2 10GE PR IOAs. You must install a combination in either of the turbo slots (slot 2 or slot 4). The 100 Gbps allocates 10 Gbps of overall bandwidth to each of these slots. With a 120 Gbps fabric configuration, the E120 router can accommodate up to 6 combinations of ES2 10G Uplink LMs and ES2-S2 10GE PR IOAs. You can install a combination in any of the line module slots, each of which are allocated 10 Gbps of overall bandwidth.

With a 320 Gbps fabric configuration, the E320 router can accommodate up to 12 combinations of ES2 10G Uplink LMs and ES2-S2 10GE PR IOAs. You can install a combination in any of the line module slots, each of which are allocated 10 Gbps of overall bandwidth.

Uplink Operation

The ES2 10G Uplink LM and ES2-S2 10GE PR IOA combination provides an uplink connection from the core network to the edge rather than traditional broadband access services such as PPPoE, transparent bridging, and subscriber interfaces.

The combination can receive and transmit full duplex line rate of 10 GB (10 GB at the ingress and 10 GB at the egress). The IOA can also support 9.6 Kbps jumbo packets at both the ingress and egress.

Multicast

The ES2 10G Uplink LM can receive multicast traffic, including all multicast control protocols. The ES2 10G Uplink LM can also transmit multicast control protocol frames and multicast data frames to perform multicast egress elaboration.

L2TP

An E-series router can be configured as either an L2TP access concentrator (LAC) or an L2TP network server (LNS). The ES2 10G Uplink LM and ES2-S2 10GE PR IOA combination supports an E-series router configured as a LAC only for traffic to or from an LNS. The ES2 10G Uplink LM and ES2-S2 10GE PR IOA combination supports both sides of the L2TP LNS function (LAC facing and core facing).

Flow Control and Policy

The ES2 10G Uplink LM and ES2-S2 10GE PR IOA combination does not support quality of service (QoS) functionality that is available on other ASIC-based Ethernet modules.

Although the ES2 10G Uplink LM does not support scheduling and shaping for egress traffic, the LM does account for the traffic class of packets through the fabric so that high priority packets are scheduled for transmission to the line module before lower priority packets. Packets that arrive at the line module are processed and transmitted using a flow-through scheme.

Currently, flow control using MAC pause frames is disabled on the ES2-S2 10GE PR IOA. The IOA does not transmit or receive pause frames. Instead, the system prioritizes control traffic over non-control traffic (that is, data). For a list of types of control traffic, see *High-Density Ethernet* on page 181.

For information about configuring policies on the ES2 10G Uplink LM and ES2-S2 10GE PR IOA, see *JUNOSe Policy Management Configuration Guide, Chapter 8, Policy Resources*.

ES2 10G LM Combination

With a 100 Gbps fabric configuration, the E320 router can accommodate up to 2 combinations of ES2 10G LMs and ES2-S2 10GE PR IOAs. You must install a combination in either of the turbo slots (slot 2 or slot 4). The 100 Gbps allocates 10 Gbps of overall bandwidth to each of these slots.

With a 120 Gbps fabric configuration, the E120 router can accommodate up to 6 combinations of ES2 10G LMs and ES2-S2 10GE PR IOAs. You can install a combination in any of the line module slots, each of which are allocated 10 Gbps of overall bandwidth.

With a 320 Gbps fabric configuration, the E320 router can accommodate up to 12 combinations of ES2 10G LMs and ES2-S2 10GE PR IOAs. You can install a combination in any of the line module slots, each of which are allocated 10 Gbps of overall bandwidth.

Access Operation

The ES2 10G LM and ES2-S2 10GE PR IOA combination provides traditional broadband access services such as PPPoE and subscriber interfaces. It also supports both sides of the L2TP LNS and LAC function.

The combination can receive and transmit full duplex line rate of 10 GB (10 GB at the ingress and 10 GB at the egress). The IOA can also support 9.6 Kbps jumbo packets at both the ingress and egress.

Multicast

The ES2 10G LM can receive multicast traffic, including all multicast control protocols. The ES2 10G LM can also transmit multicast control protocol frames and multicast data frames to perform multicast egress elaboration.

Flow Control and Policy

The ES2 10G LM and ES2-S2 10GE PR IOA combination supports QoS functionality that is available on other ASIC-based Ethernet modules.

Currently, flow control using MAC pause frames is disabled on the ES2-S2 10GE PR IOA. The IOA does not transmit or receive pause frames. Instead, the system prioritizes control traffic over non-control traffic (that is, data). For a list of types of control traffic, see *High-Density Ethernet* on page 181.

For information about configuring policies on the ES2 10G LM and ES2-S2 10GE PR IOA, see *JUNOSe Policy Management Configuration Guide, Chapter 1, Managing Policies on the E-series Router*.

Ethernet References

For more information about Ethernet implementations, consult the following resources:

- IEEE 802.1q (Virtual LANs)
- IEEE 802.1w (Rapid Reconfiguration of Spanning Tree)
- IEEE 802.3 (Fast Ethernet and Gigabit Ethernet)
- IEEE 802.3u (Fast Ethernet only)
- IEEE 802.3z (Gigabit Ethernet only)
- IEEE 802.3ae (10-Gigabit Ethernet only)
- IEEE 802.3ad (Link Aggregation)
- RFC 826—An Ethernet Address Resolution Protocol (November 1982)
- RFC 894—A Standard for the Transmission of IP Datagrams over Ethernet Networks (April 1984)
- RFC 1042—A Standard for the Transmission of IP Datagrams over IEEE 802 Networks (February 1988)
- RFC 1112—Host Extensions for IP Multicasting (August 1989)
- RFC 2516—Method for Transmitting PPP over Ethernet (PPPoE) (February 1998)

For more information about MIB support for Ethernet interfaces, consult the following resources:

- RFC 2863—The Interfaces Group MIB (June 2000)
- RFC 2668—Definitions of Managed Objects for IEEE 802.3 Medium Attachment Units (MAUs) (August 1999)
- RFC 2665—Definitions of Managed Objects for the Ethernet-like Interface Types (August 1998)

High-Density Ethernet

The following modules support high-density Ethernet:

- GE-HDE line module and GE-8 I/O module combination
- ES2 4G LM and ES2-S1 10GE IOA module combination
- ES2 4G LM and ES2-S1 GE-8 IOA module combination

In the current release, you cannot configure port parameters for high-density Ethernet. Instead, JUNOSe contains a packet classifier that enables the module to *intelligently drop* certain packets when the module becomes oversubscribed. The packet classifier inspects each incoming packet to determine whether to classify it as control traffic. To enhance network stability, the packet classifier prioritizes control traffic over non-control traffic (that is, data). The packet classifier randomly drops non-control packets when the interface is oversubscribed.

When the I/O module or IOA is oversubscribed, the packet classifier prioritizes the following types of control traffic:

- PPP discovery or PPP session
- Address Resolution Protocol (ARP)
- 802.3ad (link aggregation)
- 802.3 Spanning Tree Protocol (STP)
- IPv4 and IPv6 DHCP server
- IPv4 and IPv6 DHCP host
- IPv6 Neighbor Discovery
- IPv4 virtual router alert
- IPv4 and IPv6 Internet Group Management Protocol (IGMP)
- IPv4 packets with a type of service (ToS) precedence value set to Internetwork Control (C0)
- IPv6 packets with a traffic class precedence value set to Internetwork Control (C0)

Managing Port Redundancy on Gigabit Ethernet I/O Modules

By default, the software manages port redundancy on GE I/O modules automatically. However, you can manage redundancy on GE I/O modules as follows:

- Specify the time that the router waits for a port on a GE I/O module to become active before the router switches to the redundant port.

- Force a GE I/O module to switch operation from one port to the other.
- Disable port redundancy by specifying operation on one port only.

If you manage port redundancy manually, the router retains the manual configuration after the module reboots.

You can monitor the port redundancy configuration with the **show interfaces gigabitEthernet** command.

NOTE: The router manages failover in the same way for the GE I/O Modules and the GE-2 SFP I/O module.

link failover force

- Use to force a GE I/O module to switch operation from one port to the other.
- Select an interface on the GE I/O module before you issue this command.
- Example

  host1(config)#**interface gigabitEthernet 5/0**
  host1(config-if)#**link failover force**

- There is no **no** version.

link failover timeout

- Use to specify the time that the router waits for a port on a GE I/O module to become active before the router switches to the redundant port.
- Select an interface on the GE I/O module before you issue this command.
- Specify a time in the range 100–10,000 ms.
- Example

  host1(config)#**interface gigabitEthernet 5/0**
  host1(config-if)#**link failover timeout 1000**

- Use the **no** version to restore the default situation in which the router sets this time automatically.

link selection

- Use to disable redundancy on a GE I/O module by allowing operation on the specified port only.
- Select an interface on the GE I/O module before you issue this command.
- Example

  host1(config)#**interface gigabitEthernet 5/0**
  host1(config-if)#**link selection secondary**

- Use the **no** version to restore the default situation in which port redundancy is enabled.

Configuration Tasks for Ethernet

This section describes the options for configuring Ethernet interfaces.

You configure an Ethernet interface based on the requirements for your router configuration and the protocols you plan to route on the interface. Because you can configure an interface in different ways, Ethernet configuration tasks are divided into three primary areas. These areas are further described in separate sections in this chapter.

- Configuring the physical interface—You must perform basic configuration steps for all interfaces. This task begins with selecting an Ethernet interface and setting parameters such as line speed and MTU.
- Configuring VLANs and stacked VLANs (S-VLANs)—After you configure the physical interface, you must decide whether to configure the Ethernet interface with or without VLANs or S-VLANs. VLANs and S-VLANs enable you to multiplex multiple IP interfaces and PPPoE interfaces over a single physical Ethernet port. If you are not configuring with VLANs or S-VLANs, proceed to *Configuring Higher-Level Protocols over Ethernet* on page 225.
- Configuring higher-level protocols—You must determine which higher-level protocols, such as MPLS, to configure on the interface. This section focuses on non-VLAN configurations. You can configure some higher-level protocols, such as PPPoE, with or without VLANs.

Configuring the Physical Interface

This section describes how to complete the basic configuration for a Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet interface. CLI examples are provided with the individual command descriptions.

To configure an Ethernet interface:

1. Select an Ethernet interface.
2. (Optional) Specify the line speed and duplex mode.
3. (Optional) Specify the MTU.
4. (Optional) Set the time interval at which the router records bit and packet rates.
5. (Optional) Associate a name with the interface.
6. (Optional) Validate MAC addresses on a per interface basis.

duplex

- Use to specify the duplex mode.
- This command also works on the Fast Ethernet port on the SRP module on all E-series routers. For more information, see the *Managing the Ethernet Port on the SRP Module* in *JUNOSe System Basics Configuration Guide, Chapter 6, Managing Modules*.

NOTE: If you set either the line speed or duplex mode to automatically negotiate by using the **automatically negotiate** keyword, the router negotiates both parameters. You can specify different values to prevent the router from negotiating these parameters.

Automatic negotiation is not supported for the FE-8 SFP I/O module. For this I/O module, full duplex mode is the default.

- Example

  host1(config-if)#**duplex full**

- Use the **no** version to revert to the default, either automatically negotiate or full duplex (FE-8 SFP I/O module only)

ethernet description

- Use to associate a text description of up to 64 characters with an Ethernet interface.
- This command does not work for the Fast Ethernet port on the SRP module.
- The description is displayed in the output for **show configuration**, **show interfaces fastEthernet**, **show interfaces gigabitEthernet**, and **show interfaces tenGigabitEthernet** commands.
- Example

  host1(config-if)#**ethernet description abcd1234**

- Use the **no** version to remove the description from the interface.

interface fastEthernet

- Use to select a Fast Ethernet interface on a line module.
- You can also use it to select a Fast Ethernet management port on an SRP I/O module (ERX-7xx models, ERX-14xx models, and the ERX-310 router) or an SRP IOA (E120 and E320 routers). For information about managing the Fast Ethernet port on the SRP module, see *JUNOSe System Basics Configuration Guide, Chapter 6, Managing Modules*.
- Use the *slot/port* [*.subinterface*] format for Fast Ethernet interfaces on ERX-7xx models, ERX-14xx models, and the ERX-310 router. Use the *slot/adapter/port* format for the SRP IOA on the E120 or E320 routers; the port on the SRP IOA is always identified by 0.
- Example 1—Selects a Fast Ethernet interface on ERX-7xx models, ERX-14xx models, or the ERX-310 router

  host1(config)#**interface fastEthernet 1/0**

- Example 2—Selects the Fast Ethernet management port on an E320 router

  host1(config)#**interface fastEthernet 6/0/0**

- Use the **no** version to remove IP from an interface or subinterface. You must issue the **no** version from the highest level down; you cannot remove an interface or subinterface if the one above it still exists.

interface gigabitEthernet

- Use to select a Gigabit Ethernet interface.

NOTE: On the GE I/O module, you can configure only the primary port, 0. The router automatically uses the redundant port, 0R, if the primary port fails.

On the GE-2 SFP I/O module, you can configure only the primary ports, 0 and 1. The router automatically uses the corresponding redundant port, 0R or 1R, if the primary port fails.

On the OC3-2 GE APS I/O module, you can configure only port 2. Ports 0 and 1 are reserved for OC3/STM1 ATM interfaces. This I/O module does not support redundant ports in the current release.

On the ES2-S1 GE-4 IOA, you can configure all four ports.

On the ES2-S1 GE-8 IOA, you can configure all eight ports.

- Use the *slot/port* [*.subinterface*] format for Gigabit Ethernet interfaces on ERX-7xx models, ERX-14xx models, or the ERX-310 router; use the *slot/adapter/port* format for Gigabit Ethernet interfaces on the E120 and E320 routers.
- Example 1—Selects a Gigabit Ethernet interface on ERX-7xx models, ERX-14xx models, and the ERX-310 router

  host1(config)#**interface gigabitEthernet 1/0**
  host1(config)#**interface gigabitEthernet 2/1**

- Example 2—Selects a Gigabit Ethernet interface on the E320 router

  host1(config)#**interface gigabitEthernet 4/0/1**

- Use the **no** version to remove IP from an interface. You must issue the **no** version from the highest level down; you cannot remove an interface or subinterface if the one above it still exists.

interface tenGigabitEthernet

- Use to select a 10-Gigabit Ethernet interface on the E120 router or the E320 router.

NOTE: On the ES2-S2 10GE PR IOA, you can configure only the primary port, 0. The router automatically uses the redundant port, 0R, if the primary port fails.

- Use the *slot/adapter/port* format.

- Example—Selects a 10-Gigabit Ethernet interface on the ES2-S1 10GE IOA

  host1(config)#**interface tenGigabitEthernet 4/0/1**

- Use the **no** version to remove IP from an interface. You must issue the **no** version from the highest level down; you cannot remove an interface or subinterface if the one above it still exists.

ip mac-validate

- Use to enable or disable MAC address validation on a per interface basis.
- Use the **strict** keyword to prevent transmission of IP packets that do not reside in the validation table.
- Use the **loose** keyword to enable IP packets to pass through even though the packets do not have entries in the validation table. Only packets that have matching IP–MAC pair entries in the table are validated.
- The default behavior is not to perform MAC address validation.
- Example

  host1(config)#**interface gigabitEthernet 2/0**
  host1(config-if)#**ip address 4.4.4.2 255.255.255.0**
  host1(config-if)#**ip mac-validate strict**
  host1(config-if)#**exit**

- Use the **no** version to disable the command.

NOTE: For additional information about MAC address validation, see the **arp validate** command description in *JUNOSe IP, IPv6, and IGP Configuration Guide, Chapter 1, Configuring IP*.

load-interval

- Use to set the time interval at which the router calculates bit and packet rate counters.
- This command does not work for the Fast Ethernet port on the SRP module.
- Specify a multiple of 30 seconds, in the range 30–300 seconds.
- The default value is 300 seconds.
- Example

  host1(config-if)#**load-interval 90**

- Use the **no** version to restore the default time interval, 300 seconds.

mtu

- Use to specify the MTU for an interface.
- Specify a value in the range 64–9188 bytes. The range for FE-8 I/O modules is 64–9042 bytes.
- This command does not work for the Fast Ethernet port on the SRP module.

- Example

  host1(config-if)#**mtu 9000**

- Use the **no** version to specify the default, 1518.

speed

- Use to specify the line speed.
- This command also works on the Fast Ethernet port on the SRP module on all E-series routers. For more information, see *Managing the Ethernet Port on the SRP Module* in *JUNOSe System Basics Configuration Guide, Chapter 6, Managing Modules*.

NOTE: If you set either the line speed or duplex mode to automatically negotiate by using the **automatically negotiate** keyword, the router negotiates both parameters. You can specify different values to prevent the router from negotiating these parameters.

Automatic negotiation is not supported for the FE-8 SFP I/O module. For this I/O module, the default speed is 100 Mbps.

- Example

  host1(config-if)#**speed 10**

- Use the **no** version to revert to the default, either automatically negotiate or 100 Mbps (FE-8 SFP I/O module only).

Configuring VLANs

A virtual LAN (VLAN) enables multiplexing multiple IP and PPPoE interfaces and MPLS interfaces over a single physical Ethernet port. This multiplexing is accomplished through VLAN subinterfaces. Ethernet interfaces support the 802.1q-1998 IEEE Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks, which the router uses as its standardized format for frame tagging.

The Ethernet V2 frame format enables multiplexing of different protocols over a single physical link. IEEE 802.1q compatibility extends the frame format by adding a tag that contains a VLAN ID. This feature enables multiplexing of different channels (VLANs) over the physical link; each channel is able to multiplex different protocols.

This capability works very much like ATM encapsulation as described in RFC 2684—Multiprotocol Encapsulation over ATM Adaptation Layer 5 (September 1999). This encapsulation type enables multiplexing of multiple protocols over a single ATM virtual circuit (VC).

As shown in Figure 13, VLANs are similar to ATM VCs, with the VLAN ID serving the same function as the virtual path identifier (VPI) and virtual channel identifier (VCI) to multiplex the different channels over the physical link. The Ethernet protocol type serves the same function within a VLAN as the logical link control (LLC) subnetwork attachment point (SNAP) within a VC, to multiplex the different protocols over the channel.

Figure 13: Use of VLANs to Multiplex Different Protocols over a Single Physical Link

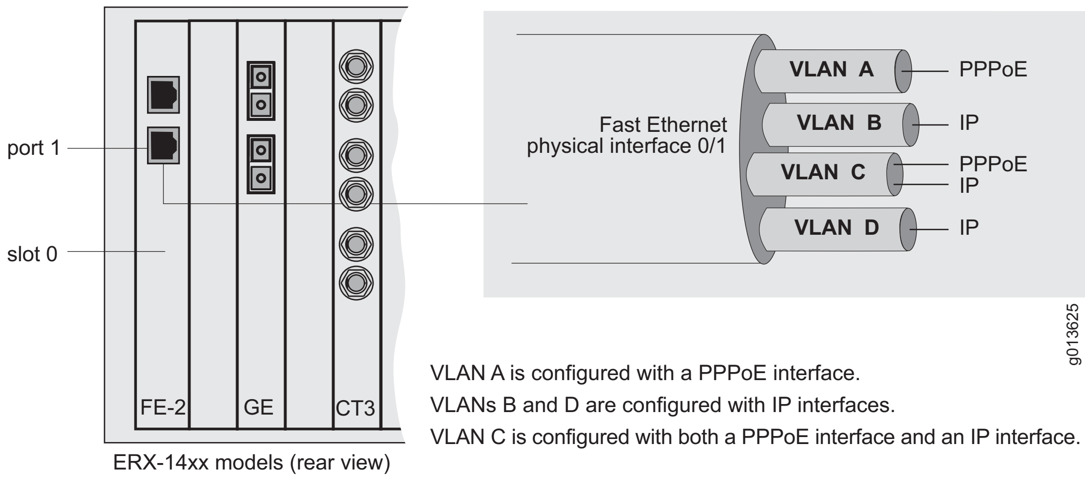

In a VLAN configuration, the router can send VLAN 0 *tagged* or *untagged frames*.

All VLAN subinterfaces use the MAC address of the Ethernet interface over which they are configured. However, some configurations, such as multiple IP over VLAN subinterfaces, require that you connect many VLAN subinterfaces to a single device. In these cases, the device uses the MAC address to identify and select the correct VLAN to use. When the MAC address is the same for all VLANs, uneven load balancing of traffic occurs. To ensure proper load balancing, you must assign unique MAC addresses to the individual VLAN subinterfaces that are connected to the device. Any ARP requests and responses generated for the IP address assigned to a VLAN subinterface use this MAC address.

You must assign the MAC address when you configure the VLAN ID. If you change the MAC address of the VLAN subinterface after you configure it, system errors can occur. To change the MAC address, you must first remove the VLAN subinterface and then reconfigure it.

For more information, see:

- *JUNOSe IP, IPv6, and IGP Configuration Guide, Chapter 1, Configuring IP*
- *JUNOSe Link Layer Configuration Guide, Chapter 7, Configuring Point-to-Point Protocol over Ethernet*

Creating a VLAN Major Interface

To use VLANs, you must first configure the Ethernet interface for VLAN encapsulation. This creates the VLAN major interface. For example:

1. Specify a Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet port.

   host1(config)#**interface fastEthernet 4/0**

2. Specify VLAN as the encapsulation method.

   host1(config-if)#**encapsulation vlan**

   The router creates the VLAN major interface.

You can now create multiple VLAN subinterfaces to carry higher-level protocols. For examples, see *Common VLAN Configurations*, next.

Common VLAN Configurations

Ethernet interfaces support IP, PPPoE, MPLS, or both IP and PPPoE on each VLAN. In addition to a VLAN major interface level, a VLAN subinterface level distinguishes the VLAN.

This section describes how to create the following common VLAN configurations, which you can configure on Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces:

- IP over VLAN
- PPPoE over VLAN
- MPLS over VLAN
- IP over VLAN and PPPoE over VLAN

NOTE: You cannot configure VLANs on the Fast Ethernet port of the SRP module.

Configuring IP over VLAN

To configure IP over VLAN over an Ethernet interface:

1. Specify a Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet port.

   host1(config)#**interface fastEthernet 4/0**

2. Specify VLAN as the encapsulation method.

   host1(config-if)#**encapsulation vlan**

   The VLAN major interface is added.

3. Create a VLAN subinterface by adding a subinterface number to the interface identification command.

   host1(config-if)#**interface fastEthernet 4/0.3**

4. Do one of the following:

   a. Assign a VLAN ID for the subinterface.

      host1(config-if)#**vlan id 201**

   b. Assign a VLAN ID and the optional unique MAC address for the subinterface.

      host1(config-if)#**vlan id 201 mac-address 0090.1a01.1234**

5. Assign an IP address and mask.

   host1(config-if)#**ip address 192.6.129.5 255.255.255.0**

6. (Optional) Configure additional VLAN subinterfaces by completing Steps 3 through 5.

Figure 14 illustrates the IP/VLAN/Fast Ethernet stacking, showing two separate VLAN subinterfaces. Configure one VLAN subinterface entirely; then configure the next VLAN subinterface.

Figure 14: Example of IP/VLAN/Fast Ethernet Stacking Configuration Steps

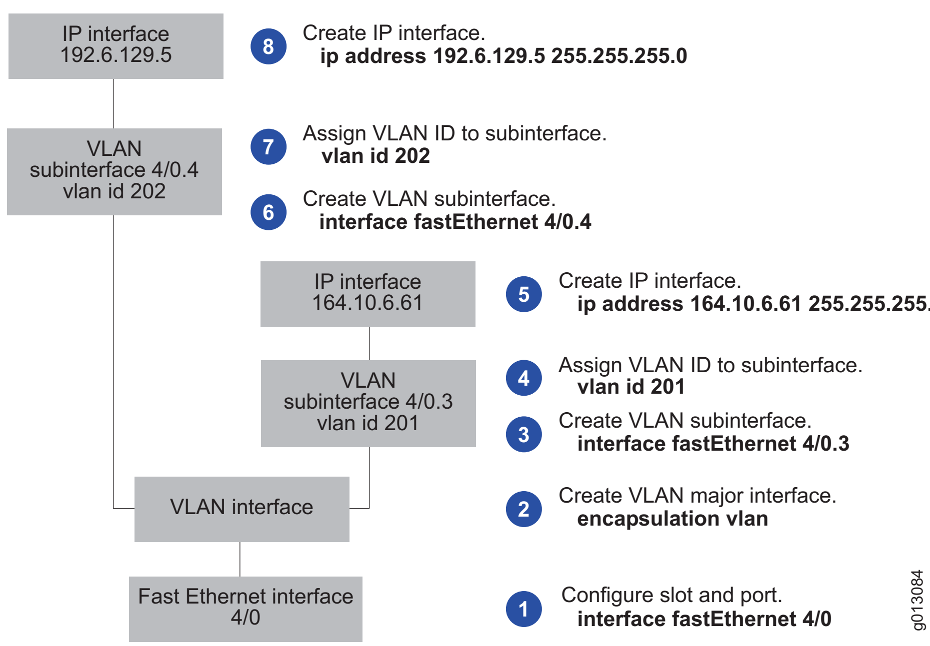

Configuring PPPoE over VLAN

To configure PPPoE over VLAN over an Ethernet interface:

1. Specify a Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet port.

   host1(config)#**interface fastEthernet 4/1**

2. Specify VLAN as the encapsulation method.

   host1(config-if)#**encapsulation vlan**

   The VLAN major interface is added.

3. Create a VLAN subinterface by adding a subinterface number to the interface identification command.

   host1(config-if)#**interface fastEthernet 4/1.1**

4. Do one of the following:

   - Assign a VLAN ID for the subinterface.

     host1(config-if)#**vlan id 201**

   - Assign a VLAN ID and the optional unique MAC address for the subinterface.

     host1(config-if)#**vlan id 201 mac-address 0090.1a01.1234**

5. Specify PPPoE as the encapsulation method on the interface.

   host1(config-if)#**pppoe**

6. Create a PPPoE subinterface.

   host1(config-if)#**pppoe subinterface fastEthernet 4/1.1.1**

7. Specify PPP as the encapsulation method on the interface.

   host1(config-if)#**encapsulation ppp**

8. Assign an IP address and mask.

   host1(config-if)#**ip address 192.6.129.5 255.255.255.0**

9. (Optional) Configure additional VLAN subinterfaces by completing Steps 3 through 8.

Figure 15 on page 192 illustrates the PPPoE/VLAN/Fast Ethernet stacking, showing two separate VLAN subinterfaces. One VLAN subinterface has two PPPoE subinterfaces, and one VLAN subinterface has one PPPoE subinterface.

Figure 15: Example of PPPoE/VLAN/Fast Ethernet Stacking Configuration Steps

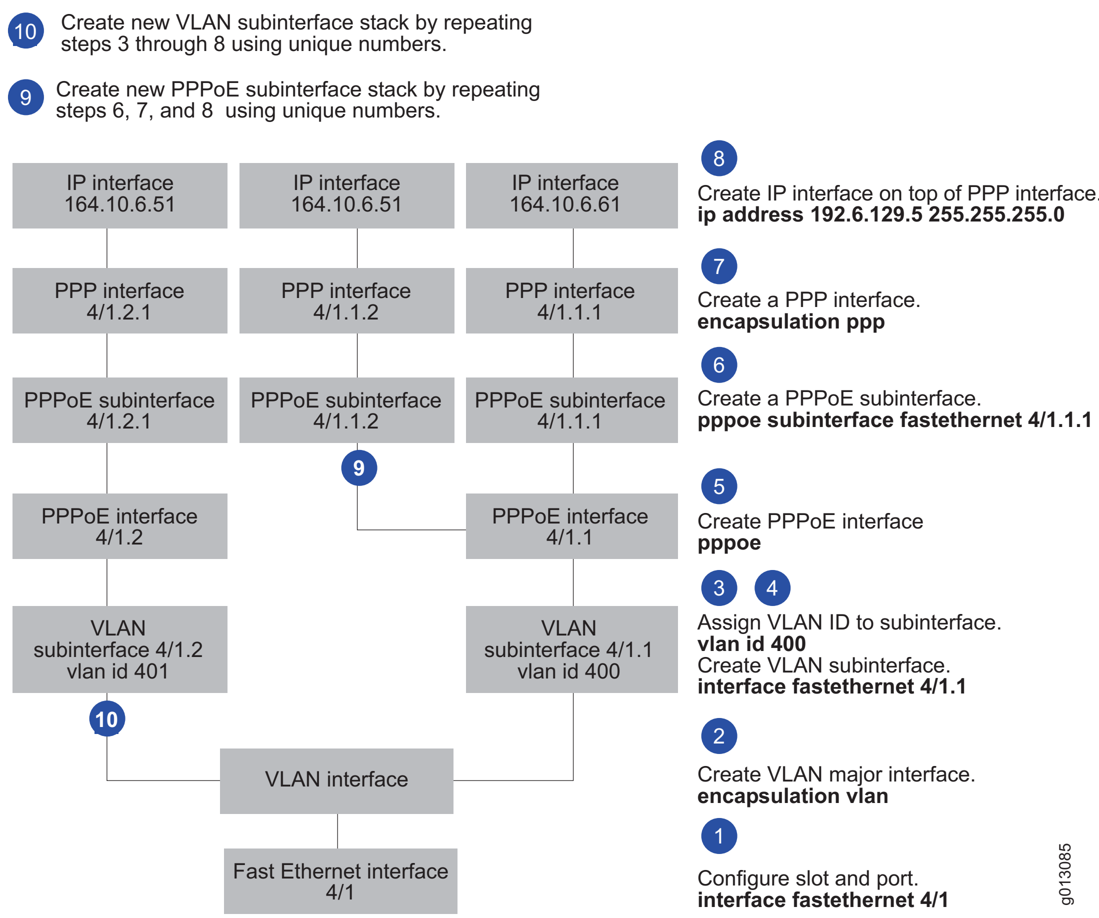

Configuring MPLS over VLAN

To configure MPLS over VLAN over an Ethernet interface:

1. Specify a Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet port.

   host1(config)#**interface fastEthernet 4/0**

2. Specify VLAN as the encapsulation method.

   host1(config-if)#**encapsulation vlan**

   The VLAN major interface is added.

3. Create a VLAN subinterface by adding a subinterface number to the interface identification command.

   host1(config-if)#**interface fastEthernet 4/1.1**

4. Do one of the following:

   - Assign a VLAN ID for the subinterface.

     host1(config-if)#**vlan id 400**

   - Assign a VLAN ID and the optional unique MAC address for the subinterface.

     host1(config-if)#**vlan id 400 mac-address 0090.1a01.1234**

5. Enable MPLS on the interface.

   host1(config-if)#**mpls**

Figure 16 illustrates the MPLS/VLAN/Fast Ethernet stacking, showing one VLAN subinterface.

Figure 16: Example of MPLS/VLAN/Fast Ethernet Stacking Configuration Steps

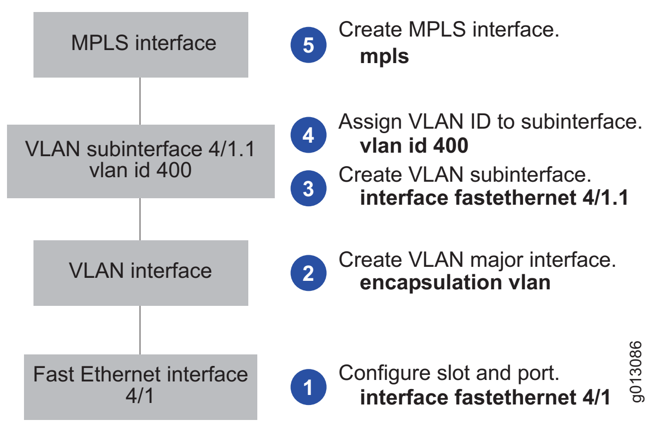

Configuring IP over VLAN and PPPoE over VLAN

To configure IP over VLAN with PPPoE over the same VLAN over an Ethernet interface:

1. Specify a Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet port.

   host1(config)#**interface fastEthernet 4/1**

2. Specify VLAN as the encapsulation method.

   host1(config-if)#**encapsulation vlan**

   The VLAN major interface is added.

3. Create a VLAN subinterface by adding a subinterface number to the interface identification command.

   host1(config-if)#**interface fastEthernet 4/1.1**

4. Do one of the following:

   - Assign a VLAN ID for the subinterface.

     host1(config-if)#**vlan id 400**

   - Assign a VLAN ID and the optional unique MAC address for the subinterface.

     host1(config-if)#**vlan id 400 mac-address 0090.1a01.1234**

5. Create an IP interface on the same VLAN as the PPPoE interface.

   host1(config-if)#**ip address 164.10.6.71 255.255.255.0**

6. Specify PPPoE as the encapsulation method on the interface.

   host1(config-if)#**pppoe**

7. Create a PPPoE subinterface.

   host1(config-if)#**pppoe subinterface fastEthernet 4/1.1.1**

8. Specify PPP as the encapsulation method on the interface.

   host1(config-if)#**encapsulation ppp**

9. Assign an IP address and mask.

   host1(config-if)#**ip address 192.6.129.5 255.255.255.0**

10. (Optional) Configure additional PPPoE subinterfaces by completing Steps 7 through 9 using unique numbering.

To configure additional IP interfaces over the VLAN major interface:

1. Create a new VLAN subinterface by adding a unique subinterface number to the interface identification command.

   host1(config-if)#**interface fastEthernet 4/1.2**

2. Assign a VLAN ID for the subinterface.

   host1(config-if)#**vlan id 401**

3. Assign an IP address and mask.

   host1(config-if)#**ip address 164.10.6.51 255.255.255.0**

Figure 17 illustrates the configuration steps for two VLAN subinterfaces. In this example:

- VLAN subinterface 4/1.1 has an IP interface, a PPPoE interface, and multiple PPPoE subinterface stacks.
- VLAN subinterface 4/1.2 has only an IP interface.

NOTE: Before you can remove a VLAN subinterface, you must remove the upper-layer interface stack.

Figure 17: Example of PPPoE over VLAN with IP over VLAN Stacking Configuration Steps

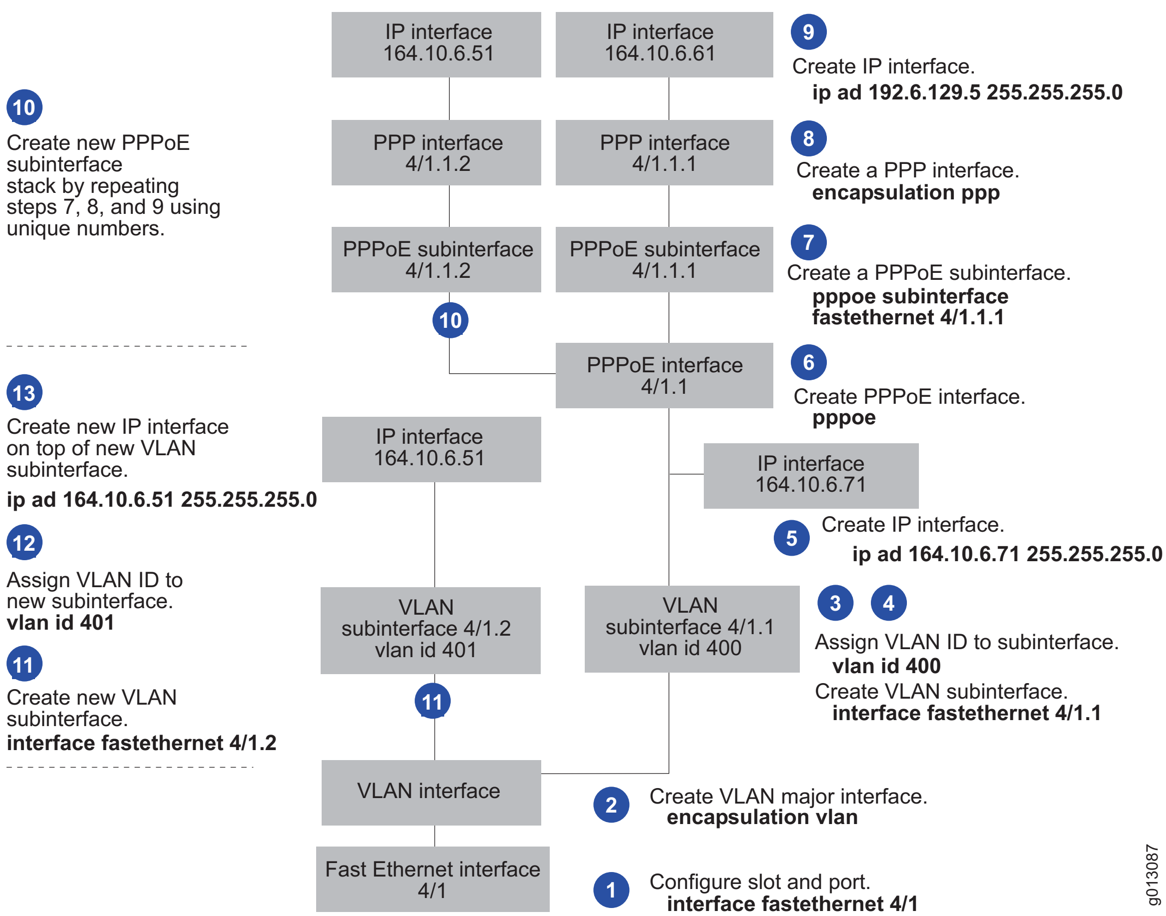

encapsulation ppp

- Use to configure PPP as the encapsulation method for the interface.
- Example

  host1(config-if)#**encapsulation ppp**

- Use the **no** version to disable PPP on the interface.

encapsulation vlan

- Use to configure VLAN as the encapsulation method for the interface.
- Example

  host1(config-if)#**encapsulation vlan**

- Use the **no** version to disable VLAN on an interface.

ip address

- Use to set a primary or secondary IP address for an interface or subinterface.
- Specify the layer 2 encapsulation before you set the IP address.
- Example

  host1(config-if)#**ip address 192.6.129.5 255.255.255.0**

- Use the **no** version to remove an IP address or disable IP processing.

pppoe

- Use to configure PPPoE as the encapsulation method on the interface.
- Example

  host1(config-if)#**pppoe**

- Use the **no** version to disable PPPoE on the interface.

pppoe subinterface fastEthernet

- Use to create a PPPoE subinterface on a Fast Ethernet interface.
- Example

  host1(config-if)#**pppoe subinterface fastEthernet 4/1.1.1**

- Use the **no** version to remove a PPPoE subinterface on a Fast Ethernet interface.

pppoe subinterface gigabitEthernet
pppoe subinterface tenGigabitEthernet

- Use to create a PPPoE subinterface on a Gigabit Ethernet interface or on a 10-Gigabit Ethernet interface.
- Example 1—Creates a PPPoE subinterface on an ERX-7xx model, ERX-14xx model, or the ERX-310 router

  host1(config-if)#**pppoe subinterface gigabitEthernet 4/2.1.1**

- Example 2—Creates a PPPoE subinterface on the E320 router

  host1(config-if)#**pppoe subinterface tenGigabitEthernet 4/0/2.1.1**

- Use the **no** version to remove a PPPoE subinterface on a Gigabit Ethernet interface or on a 10-Gigabit Ethernet interface.

vlan description

- Use to assign an alias or description to a VLAN subinterface.
- You can use a maximum of 64 characters for the description or to name the alias.
- Example

  host1(config-if)#**vlan description randolph56a**

- Use the **no** version to remove the VLAN description.

vlan id

- Use to specify the VLAN ID.
- Use a VLAN ID that is in the range 0–4095 and is unique within the Ethernet interface.
- Issue the **vlan id** command before any upper bindings are made, such as IP or PPPoE.
- Use the **mac-address** keyword to specify a unique MAC address for the VLAN subinterface. When you do not specify a unique MAC address, the VLAN uses the MAC address of the Ethernet interface.
- Use the optional keyword **untagged** to specify that frames be sent untagged. The keyword is valid only for VLAN ID 0. Tagged frames can be received, but untagged frames are sent.
- Examples

  host1(config-if)#**vlan id 400**
  host1(config-if)#**vlan id 4 255 mac-address 0090.1a01.1234**

- There is no **no** version.

Configuring S-VLANs

As described in *Configuring VLANs* on page 187, VLANs permit multiplexing multiple IP interfaces and PPPoE interfaces over a single physical Ethernet port by creating VLAN subinterfaces. As specified in IEEE Standard 802.1q, the 12-bit VLAN identifier's tagged frames enables the construction of a maximum of 4096 distinct VLANs. In an Ethernet B-RAS application environment, however, this VLAN limit is inadequate. A stacked VLAN (S-VLAN) provides a two-level VLAN tag structure, which extends the VLAN ID space to more than 16 million VLANs.

Creating an S-VLAN requires the use of a second encapsulation tag. The router performs decapsulation twice, once to get the S-VLAN tag and once to get the VLAN tag. This *double tagging* approach enables more than 16 million address possibilities, which more than satisfies the scaling requirement for Ethernet B-RAS applications.

VLAN and S-VLAN subinterfaces can coexist over the same VLAN major interface. You configure S-VLANs in the same way that you configure VLANs, with the addition of certain commands.

NOTE: See *JUNOSe Release Notes*, *Appendix A*, *System Maximums* for S-VLAN limitations.

Like VLANs, all S-VLAN subinterfaces use the MAC address of the Ethernet interface over which they are configured. For more information about assigning unique MAC address to the S-VLAN subinterface when assigning VLAN or S-VLAN IDs, see *Configuring VLANs* on page 187.

Configuring PPPoE over S-VLAN

To configure PPPoE over an S-VLAN over an Ethernet interface:

1. Specify a Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet port.

   host1(config)#**interface fastEthernet 4/0**

2. Specify VLAN as the encapsulation method.

   host1(config-if)#**encapsulation vlan**

   The VLAN major interface is added.

3. Create a VLAN subinterface by adding a subinterface number to the interface identification command.

   host1(config-if)#**interface fastEthernet 4/1.1**

4. Assign an S-VLAN ID and a VLAN ID for the subinterface.

   host1(config-if)#**svlan id 4 255**

5. Assign an S-VLAN Ethertype.

   host1(config-if)#**svlan ethertype 88a8**

6. Specify PPPoE as the encapsulation method on the interface.

   host1(config-if)#**pppoe**

7. Create a PPPoE subinterface.

   host1(config-if)#**pppoe subinterface fastEthernet 4/1.1.1**

8. Specify PPP as the encapsulation method on the interface.

   host1(config-if)#**encapsulation ppp**

9. Assign an IP address and mask.

   host1(config-if)#**ip address 164.10.6.61 255.255.255.0**

10. (Optional) Configure additional PPPoE subinterfaces by completing Steps 7 through 9 using unique numbering.

Figure 18 shows one S-VLAN subinterface with multiple PPPoE subinterface stacks.

NOTE: Before you can remove an S-VLAN/VLAN subinterface, you must remove the upper-layer interface stack.

Figure 18: Example of PPPoE over S-VLAN Stacking Configuration Steps

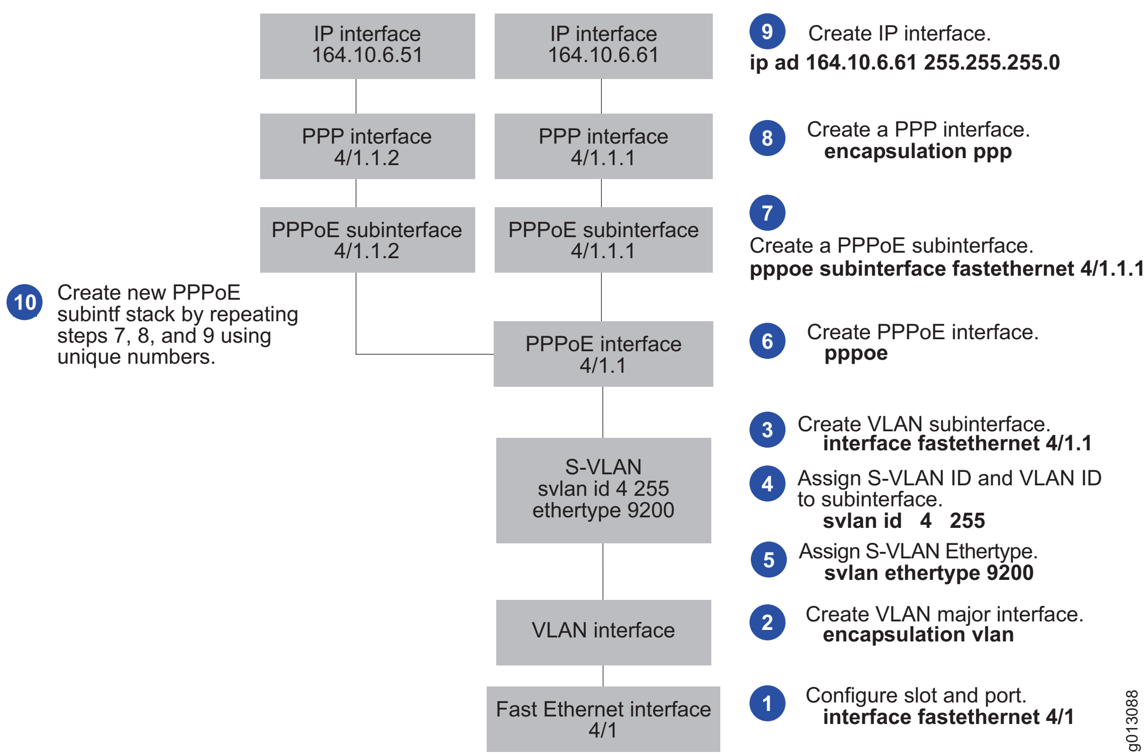

encapsulation ppp

- Use to configure PPP as the encapsulation method for the interface.
- Use the **no** version to remove PPP as the encapsulation method on the interface.

encapsulation vlan

- Use to configure VLAN as the encapsulation method for the interface.
- Use the **no** version to remove VLAN as the encapsulation method on the interface.

ip address

- Use to set a primary or secondary IP address for an interface or subinterface.
- Specify the layer 2 encapsulation before you set the IP address.
- Use the **no** version to remove an IP address or disable IP processing.

pppoe

- Use to configure PPPoE as the encapsulation method on the interface.
- Use the **no** version to disable PPPoE on the interface.

pppoe subinterface fastEthernet

- Use to create a PPPoE subinterface on a Fast Ethernet interface.
- Use the **no** version to remove a PPPoE subinterface on a Fast Ethernet interface.

pppoe subinterface gigabitEthernet
pppoe subinterface tenGigabitEthernet

- Use to create a PPPoE subinterface on a Gigabit Ethernet interface or on a 10-Gigabit Ethernet interface.
- Use the **no** version to remove a PPPoE subinterface on a Gigabit Ethernet interface or on a 10-Gigabit Ethernet interface.

svlan ethertype

- Use to assign an Ethertype value for the S-VLAN subinterface.
- Choose one of the following Ethertype values:
  - 8100—Specifies Ethertype value 0x8100, as defined in IEEE Standard 802.1q
  - 88a8—Specifies Ethertype value 0x88a8, as defined in draft IEEE Standard 802.1ad
  - 9100—Specifies Ethertype value 0x9100, which is the default
- Use an Ethertype value that matches the Ethertype value set on the customer premises equipment (CPE) to which your router connects.
- Example

  host1(config-if)#**svlan ethertype 8100**

- Use the **no** version to restore the default value, 9100.

svlan id

- Use to assign S-VLAN IDs and VLAN IDs to VLAN subinterfaces.
- Use S-VLAN ID and VLAN ID numbers that are in the range 0–4095 and that are unique within the Ethernet interface.
- Use the **mac-address** keyword to specify a unique MAC address for the VLAN subinterface. When you do not specify a unique MAC address, the VLAN uses the MAC address of the Ethernet interface.
- Examples

  host1(config-if)#**svlan id 4 255**
  host1(config-if)#**svlan id 4 255 mac-address 0090.1a01.1234**

- Issue the **svlan id** command before any upper bindings are made, such as IP or PPPoE.
- There is no **no** version.

Configuring S-VLAN Tunnels for Layer 2 Services over MPLS

When you configure Ethernet layer 2 services over MPLS, you can create a special type of S-VLAN called an S-VLAN tunnel that uses a single interface to tunnel traffic from multiple VLANs across an MPLS network. The S-VLAN tunnel enables multiple VLANs, each configured with a unique VLAN ID tag, to share a common S-VLAN ID tag when they traverse an MPLS network.

Advantages

Using S-VLAN tunnels provides an easier and faster way to configure Ethernet layer 2 services over MPLS than using standard S-VLANs. For example, consider the network configuration shown in Figure 19.

Figure 19: S-VLAN Tunnels for Ethernet Layer 2 Services over MPLS

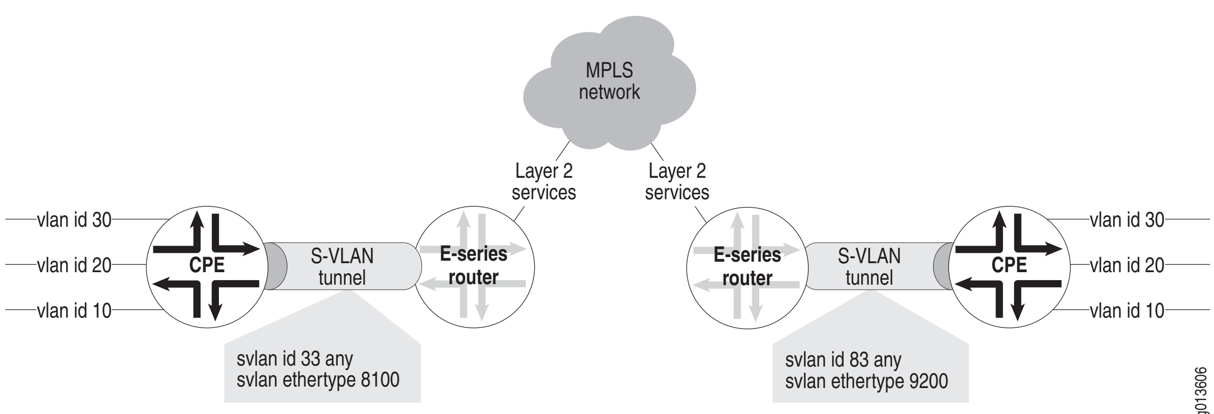

In this example, traffic from three VLAN subinterfaces must traverse the MPLS network. To accomplish this using standard S-VLANs, you issue the following commands to configure three separate S-VLANs with the same S-VLAN ID value and different VLAN IDs, as follows:

```
host1(config-if)#svlan id 33 10
host1(config-if)#svlan id 33 20
host1(config-if)#svlan id 33 30
```

By contrast, using an S-VLAN tunnel achieves the same result, but requires you to issue only a single **svlan id** command with the keyword **any** in place of the VLAN ID value. For example, the following command creates a single interface that tunnels traffic from VLANs configured with an S-VLAN ID of 33 and *any* VLAN ID to the same destination across the MPLS network. In effect, this command tunnels traffic from all three VLANs shown in Figure 19 on page 201.

```
host1(config-if)#svlan id 33 any
```

Interface Stacking

When you configure Ethernet layer 2 services over MPLS using S-VLAN tunnels, the only interface that you can stack over an S-VLAN tunnel is an MPLS tunnel, which you configure using the MPLS tunneling command (**mpls-relay** or **route interface**) that is appropriate for your configuration. Attempting to configure any other interface type—such as IP, MPLS (nontunnel), or PPPoE—over the S-VLAN tunnel causes the router to generate an error and reject the configuration as invalid.

For details about configuring MPLS and layer 2 services over MPLS, see:

- *JUNOSe BGP and MPLS Configuration Guide, Chapter 2, Configuring MPLS*
- *JUNOSe BGP and MPLS Configuration Guide, Chapter 5, Configuring Layer 2 Services over MPLS*

Configuration Example

This section uses the sample network topology shown in Figure 19 on page 201 to illustrate the steps for configuring S-VLAN tunnels for Ethernet layer 2 services over MPLS.

To configure S-VLAN tunnels for Ethernet layer 2 services over MPLS:

1. Specify a Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet port.

   ```
   host1(config)#interface fastEthernet 4/0
   ```

2. Specify VLAN as the encapsulation method to create the VLAN major interface.

   ```
   host1(config-if)#encapsulation vlan
   ```

3. Create a VLAN subinterface.

   ```
   host1(config-if)#interface fastEthernet 8/1.1
   ```

4. Create the S-VLAN tunnel. This interface tunnels traffic from VLANs configured with an S-VLAN ID of 33 and any VLAN ID to the same destination across the MPLS network.

   host1(config-if)#**svlan id 33 any**

5. Assign an S-VLAN Ethertype.

   host1(config-if)#**svlan ethertype 8100**

6. Create the MPLS tunnel interface using the appropriate MPLS tunneling command for your configuration. For example:

   host1(config-if)#**route interface tunnel mpls:tunnel3 45**

   For complete instructions on configuring the MPLS tunnel, see *JUNOSe BGP and MPLS Configuration Guide, Chapter 5, Configuring Layer 2 Services over MPLS*.

7. Repeat Steps 1 through 6 using unique values to configure the S-VLAN tunnel and MPLS tunnel interfaces on the remote E-series router. For example:

   host2(config)#**interface fastEthernet 3/1**
   host2(config-if)#**encapsulation vlan**
   host2(config-if)#**interface fastEthernet 3/1.1**
   host2(config-if)#**svlan id 83 any**
   host2(config-if)#**svlan ethertype 88a8**
   host2(config-if)#**route interface tunnel mpls:tunnel2 45**

encapsulation vlan

- Use to configure VLAN as the encapsulation method for the interface.
- Use the **no** version to disable VLAN on an interface.

interface fastEthernet

- Use to select a Fast Ethernet interface on a line module.
- Example

  host1(config)#**interface fastEthernet 3/1**

- Use the **no** version to remove the interface or subinterface. You must issue the **no** version from the highest level down; you cannot remove an interface or subinterface if the one above it still exists.

route interface

- Use to route layer 2 traffic on a specific tunnel interface.
- Use the **no** version to negate this command.

NOTE: For details on the use of this command, see *JUNOSe BGP and MPLS Configuration Guide, Chapter 5, Configuring Layer 2 Services over MPLS*.

svlan ethertype

- Use to assign an Ethertype value for the S-VLAN tunnel interface.
- Choose one of the following Ethertype values:
  - 8100—Specifies Ethertype value 0x8100, as defined in IEEE Standard 802.1q
  - 88a8—Specifies Ethertype value 0x88a8, as defined in draft IEEE Standard 802.1ad
  - 9100—Specifies Ethertype value 0x9100, which is the default
- Use an Ethertype value that matches the Ethertype value set on the customer premises equipment (CPE) to which your router connects.
- Example

  host1(config-if)#**svlan ethertype 8100**

- Use the **no** version to restore the default value, 9100.

svlan id

- Use to create an S-VLAN tunnel interface for configuring Ethernet layer 2 services over MPLS.
- Assign an S-VLAN ID value in the range 0–4095 that is unique within the Ethernet interface.
- Use the **any** keyword to tunnel traffic from VLANs configured with the specified S-VLAN ID and any VLAN ID to the same destination across an MPLS network.
- Issue the **svlan id** command with the **any** keyword before you configure the upper binding, which must be an MPLS tunnel interface. Attempting to configure any other interface type over the S-VLAN tunnel causes an error.
- Example

  host1(config-if)#**svlan id 1000 any**

- There is no **no** version.

S-VLAN Oversubscription

When you configure S-VLAN subinterfaces over Ethernet interfaces to support dynamic PPPoE subinterfaces, you can take advantage of S-VLAN oversubscription.

The following module combinations support S-VLAN oversubscription:

- GE/FE line module and all of its associated I/O modules
- GE-2 line module and the GE-2 SFP I/O module
- GE-HDE line module and its associated I/O modules
- OC3/STM1 GE/FE line module and the OC3-2 GE APS I/O module
- ES2 4G LM and its associated Gigabit Ethernet and 10-Gigabit Ethernet IOAs
- ES2 10G LM and its associated Gigabit Ethernet and 10-Gigabit Ethernet IOAs

The maximum number of S-VLANs that you can create per I/O module with PPPoE major interfaces stacked over them is greater than the maximum number of dynamic PPPoE subinterfaces. The maximum number of PPP interfaces supported per line module is directly proportional to the maximum number of PPPoE subinterfaces.

As a result, you can oversubscribe S-VLANs by configuring up to the maximum number of S-VLANs supported on these I/O modules, knowing that no more than the maximum number of supported PPP sessions can be connected to the router at any one time.

For configuration instructions, see *Configuring Dynamic PPPoE over Static PPPoE with Ethernet and S-VLAN Interface Columns* in *JUNOSe Link Layer Configuration Guide, Chapter 12, Configuring Dynamic Interfaces*.

For specific information about the maximum number of S-VLANs supported per I/O module and the maximum number of PPP interfaces and PPPoE subinterfaces supported per line module, see *JUNOSe Release Notes*, *Appendix A*, *System Maximums*.

NOTE: The E120 and E320 routers can support up to two IOAs per line module. This maximum number of S-VLANs per line module does not change if one or two IOAs are installed.

Configuring 802.3ad Link Aggregation for Ethernet

IEEE 802.3ad link aggregation enables you to group Ethernet interfaces at the physical layer to form a single link layer interface, also known as a link aggregation group (LAG) or bundle. For more information, see IEEE Standard 802.3ad, Link Aggregation.

Some users require more bandwidth in their network than a single Fast Ethernet link can provide, but cannot afford the expense or do not need the bandwidth of a higher-speed Gigabit Ethernet link. Using IEEE 802.3ad link aggregation in this situation provides increased port density and bandwidth at lower cost. For example, if you need 450 Mbps of bandwidth to transmit data and have only a 100-Mbps Fast Ethernet link, creating a LAG bundle containing five 100-Mbps Fast Ethernet links is more cost effective than purchasing a single Gigabit Ethernet link.

For information about the modules that support link aggregation, see *ERX Module Guide, Appendix A, Module Protocol Support* and *E120 and E320 Module Guide, Appendix A, IOA Protocol Support*.

LACP

The Link Aggregation Control Protocol (LACP) is a mechanism for exchanging port and system information to create and maintain LAG bundles. The LAG bundle distributes MAC clients across the link layer interface and collects traffic from the links to present to the MAC clients of the LAG bundle.

To create the links in the LAG bundles, you can add one or more Ethernet physical interfaces to it. The LACP detects Ethernet interfaces as links if they are configured on the same line module and have the same physical layer characteristics. The LACP also assigns to the LAG bundle the same MAC address of the Ethernet link with the highest port priority, which is the lowest value.

The LACP also controls the exchange of LACP protocol data units (PDUs) between the Ethernet links in the LAG bundle. The PDUs contain information about each link and enable the LAG bundle to maintain them.

By default, Ethernet links do not exchange PDUs, which contain information about the state of the link. You can configure Ethernet links to actively transmit PDUs, or passively transmit them, sending out LACP PDUs only when it receives them from another link. The transmitting link is known as the *Actor* and the receiving link is known as the *Partner*.

Higher-Level Protocols

After you configure the LAG bundle, you can route IP traffic over it, create a VLAN over it, route PPPoE traffic over it, or route MPLS traffic over it.

Figure 20 displays the interface stack for 802.3ad link aggregation.

Figure 20: Interface Stack for 802.3ad Link Aggregation

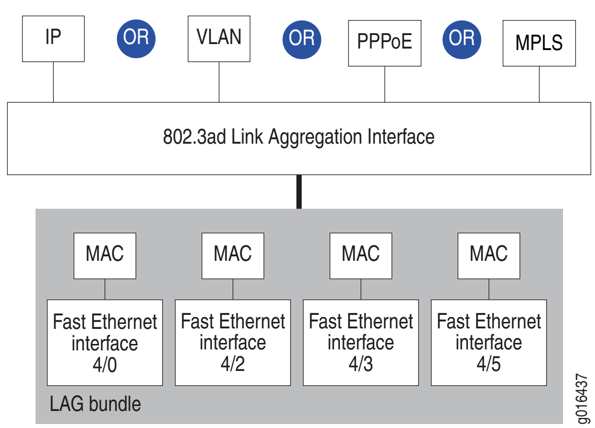

For information about configuring higher-level protocols over VLANs, see *Common VLAN Configurations* on page 189.

Load Balancing and QoS

You can configure load balancing across 802.3ad links to provide quality of service (QoS). To ensure that QoS is symmetrically applied to all the links, the router periodically rebalances the traffic on the LAG. When you attach a QoS profile to the LAG, the load balancing properties that are configured are applied to the LAG, and determines how traffic is distributed.

For example, if VLANs are configured, IP queues are provisioned over the VLANs. In this case, the default behavior is per-VLAN load balancing.

For more information, see *JUNOSe Quality of Service Configuration Guide, Chapter 20, Configuring QoS for Gigabit Ethernet Interfaces and VLAN Subinterfaces*.

Configuration Tasks for 802.3ad Link Aggregation

To configure link aggregation on Ethernet interfaces, you must configure the Ethernet interface, create the LAG bundle, and add the Ethernet interface as a member link in the LAG bundle. Optionally, you can then configure IP, a VLAN subinterface, a PPPoE subinterface, or MPLS for the LAG bundle.

For more information about specifying LAG interfaces and subinterfaces on E-series routers, see *Interface Types and Specifiers* in *JUNOSe Command Reference Guide, About This Guide*.

Configuring an Ethernet Physical Interface

To configure a member link, perform the following steps:

1. Specify a Fast Ethernet or Gigabit Ethernet interface for which you want to create a member link.

   host1(config)#**interface gigabitEthernet 2/0**

2. Configure LACP in passive or active mode.

   host1(config-if)#**lacp active**

3. Specify the speed and the duplex mode for the Ethernet interface.

   host1(config-if)#**speed 100**
   host1(config-if)#**duplex full**

4. To configure additional member links, repeat steps 1 to 3.

NOTE: All of the member links that you configure must have the same physical layer characteristics, such as speed and duplex mode.

Configuring a LAG Bundle

To configure a LAG bundle and add member links, perform the following steps:

1. Create the LAG bundle.

   host1(config)#**interface lag bundleBoston**

2. Add a member link to the LAG bundle.

   host1(config-if)#**member–interface gigabitEthernet 2/0**

Configuring IP for a LAG Bundle

To configure IP for a LAG bundle, perform the following steps:

1. Specify the LAG bundle.

   host1(config)#**interface lag bundleBoston**

2. Assign an IP address and mask.

   host1(config-if)#**ip address 192.5.127.8 255.255.255.0**

Configuring a VLAN Subinterface for a LAG Bundle

To configure a VLAN subinterface for the LAG bundle, perform the following steps:

1. Specify VLAN as the encapsulation method.

   host1(config-if)#**encapsulation vlan**

2. Specify the VLAN subinterface for the LAG bundle by adding a unique subinterface number to the LAG interface identification command.

   host1(config)# **interface lag bundleBoston.1**

3. Assign a VLAN ID for the subinterface.

   host1(config-if)#**vlan id 203**

4. Assign an IP address and mask.

   host1(config-if)#**ip address 192.168.1.1 255.255.0.0**

Configuring a PPPoE Subinterface for a LAG Bundle

To configure a PPPoE subinterface for the LAG bundle, perform the following steps:

1. Specify PPPoE as the encapsulation method.

   host1(config-if)#**encapsulation pppoe**

2. Specify the PPPoE subinterface for the LAG bundle in either of the following ways:

- Use the **interface lag** command to add a unique subinterface number to the LAG bundle name.

  host1(config)#**interface lag bundleBoston.2**

- Use the **pppoe subinterface lag** command to add a unique subinterface number to the LAG bundle name.

  host1(config)#**pppoe subinterface lag bundleBoston.2**

3. Specify PPP as the encapsulation method on the PPPoE subinterface.

   host1(config-if)#**encapsulation ppp**

4. Assign an IP address and mask.

   host1(config-if)#**ip address 192.168.1.2 255.255.0.0**

You can also configure a PPPoE subinterface over a VLAN subinterface over a LAG bundle. For an example of this configuration, see *Example: Configuring a PPPoE Subinterface over a VLAN for a LAG Bundle* on page 211.

Configuring MPLS for a LAG Bundle

To configure MPLS for a LAG bundle, perform the following steps:

1. Specify the LAG bundle.

   host1(config)#**interface lag bundleBoston**

2. Create an MPLS interface.

   host1(config-if)#**mpls**

802.3ad Link Aggregation Configuration Examples

This section provides examples for the following 802.3ad link aggregation configurations:

- IP interface over a LAG bundle
- PPPoE subinterface over a LAG bundle
- PPPoE subinterface over a VLAN subinterface over a LAG bundle
- MPLS over a LAG bundle
- MPLS over a VLAN subinterface over a LAG bundle

Example: Configuring an IP Interface for a LAG Bundle

The following example displays configuration of LACP for two Fast Ethernet interfaces in slot 0. The interfaces are enabled for active LACP. The speed and duplex characteristics are the same for both interfaces.

```
host1(config)#interface fastEthernet 0/0
host1(config-if)#speed 100
host1(config-if)#duplex full
host1(config-if)#lacp active
host1(config-if)#interface fastEthernet 0/5
host1(config-if)#speed 100
host1(config-if)#duplex full
host1(config-if)#lacp active
```

The following commands create a virtual router, add the Ethernet physical interfaces to a LAG bundle named bundleBoston, and assign an IP address and mask to the bundle.

```
host1(config)#virtual-router boston
host1:boston(config)#interface lag boston
host1:boston(config-if)#member-interface fastEthernet 0/0
host1:boston(config-if)#member-interface fastEthernet 0/5
host1:boston(config-if)#ip address 1.1.1.1 255.255.255.0
```

Example: Configuring a PPPoE Subinterface for a LAG Bundle

The following example displays LACP configuration for two Fast Ethernet interfaces in slot 4. The interfaces are enabled for passive LACP. The speed and duplex characteristics are the same for both interfaces.

```
host1(config)#interface fastEthernet 4/0
host1(config-if)#speed 100
host1(config-if)#duplex full
host1(config-if)#lacp passive
host1(config-if)#interface fastEthernet 4/3
host1(config-if)#speed 100
host1(config-if)#duplex full
host1(config-if)#lacp passive
```

The following commands add the Ethernet physical interfaces to a LAG bundle named chicago.

```
host1(config)#interface lag chicago
host1(config-if)#member-interface fastEthernet 4/0
host1(config-if)#member-interface fastEthernet 4/3
```

The following commands configure a PPPoE subinterface for the LAG bundle named chicago. In the LAG interface identification command (**interface lag chicago.1**), the number 1 represents the subinterface number for the PPPoE subinterface.

```
host1(config-if)#encapsulation pppoe
host1(config)#interface lag chicago.1
host1(config-if)#encapsulation ppp
host1(config-if)#ip address 10.10.1.1 255.255.0.0
```

As an alternative to using the command **interface lag chicago.1** to configure the PPPoE subinterface in this example, you can also use the command **pppoe subinterface lag chicago.1** to achieve the same result. For more information, see **pppoe subinterface lag** on page 214.

Example: Configuring a PPPoE Subinterface over a VLAN for a LAG Bundle

The following example displays LACP configuration for two Fast Ethernet interfaces in slot 3. The interfaces are enabled for active LACP. The speed and duplex characteristics are the same for both interfaces.

```
host1(config)#interface fastEthernet 3/0
host1(config-if)#speed 100
host1(config-if)#duplex full
host1(config-if)#lacp active
host1(config-if)#interface fastEthernet 3/1
host1(config-if)#speed 100
host1(config-if)#duplex full
host1(config-if)#lacp active
```

The following commands add the Ethernet physical interfaces to a LAG bundle named sunnyvale.

```
host1(config)#interface lag sunnyvale
host1(config-if)#member-interface fastEthernet 3/0
host1(config-if)#member-interface fastEthernet 3/1
```

The following commands configure a VLAN subinterface for the LAG bundle named sunnyvale. In the LAG interface identification command (**interface lag sunnyvale.1**), the number 1 represents the subinterface number for the VLAN subinterface.

```
host1(config-if)#encapsulation vlan
host1(config)#interface lag sunnyvale.1
host1(config-if)#vlan id 100
```

The following commands configure a PPPoE subinterface over the VLAN subinterface for the LAG bundle named sunnyvale. In the LAG interface identification command (**interface lag sunnyvale.1.2**), the number 2 represents the subinterface number for the PPPoE subinterface.

```
host1(config-if)#encapsulation pppoe
host1(config)#interface lag sunnyvale.1.2
host1(config-if)#encapsulation ppp
host1(config-if)#ip address 10.10.2.2 255.255.0.0
```

As an alternative to using the command **interface lag sunnyvale.1.2** to configure the PPPoE subinterface in this example, you can also use the command **pppoe subinterface lag sunnyvale.1.2** to achieve the same result. For more information, see **pppoe subinterface lag** on page 214.

Example: Configuring MPLS for a LAG Bundle

The following example displays configuration of LACP for two Fast Ethernet interfaces in slot 5. The interfaces are enabled for active LACP. The speed and duplex characteristics are the same for both interfaces.

```
host1(config)#interface fastEthernet 5/0
host1(config-if)#speed 100
host1(config-if)#duplex full
host1(config-if)#lacp active
host1(config-if)#interface fastEthernet 5/1
host1(config-if)#speed 100
host1(config-if)#duplex full
host1(config-if)#lacp active
```

The following commands create a virtual router, add the Ethernet physical interfaces to a LAG bundle named kanata, assign an IP address, and configure MPLS.

```
host1(config)#virtual router kanata
host1:kanata(config)#interface lag kanata
host1:kanata(config-if)#member-interface fastEthernet 0/0
host1:kanata(config-if)#member-interface fastEthernet 0/5
host1:kanata(config-if)#ip address 1.1.1.1 255.255.255.0
host1(config-if)#mpls
```

Example: Configuring MPLS over a VLAN for a LAG Bundle

The following example displays configuration of LACP for two Fast Ethernet interfaces in slot 5. The interfaces are enabled for active LACP. The speed and duplex characteristics are the same for both interfaces.

```
host1(config)#interface fastEthernet 5/0
host1(config-if)#speed 100
host1(config-if)#duplex full
host1(config-if)#lacp active
host1(config-if)#interface fastEthernet 5/1
host1(config-if)#speed 100
host1(config-if)#duplex full
host1(config-if)#lacp active
```

The following commands add the Ethernet physical interfaces to a LAG bundle named kanata.

```
host1(config)#virtual router kanata
host1:kanata(config)#interface lag kanata
host1:kanata(config-if)#member-interface fastEthernet 5/0
host1:kanata(config-if)#member-interface fastEthernet 5/1
```

The following commands configure a VLAN subinterface for the LAG bundle named kanata. In the LAG interface identification command (**interface lag kanata.1**), the number 1 represents the subinterface number for the VLAN subinterface.

host1:kanata(config-if)#**encapsulation vlan**
host1:kanata(config)#**interface lag kanata.1**
host1:kanata(config-if)#**vlan id 100**

The following command creates an MPLS interface.

host1:kanata(config)#**mpls**

interface lag

- Use to create an IEEE 802.3ad LAG interface, also known as an LAG bundle, or a subinterface for the LAG bundle.
- Examples

  host1(config)#**interface lag boston**
  host1(config)#**interface lag boston.2**
  host1(config)#**interface lag boston.2.1**

- Use the **no** version to delete the LAG bundle.

lacp

- Use to configure whether an Ethernet link in a LAG bundle participates actively or passively in the LACP.
- Use the **active** keyword to indicate that the Ethernet link participates in the protocol regardless of whether its Partner member link is set to active or passive LACP PDU participation.
- Use the **passive** keyword to indicate that the Ethernet link to transmit LACP PDUs only when it receives LACP PDUs from its Partner member link.
- By default, Ethernet links in a LAG bundle do not send LACP PDUs.
- Example

  host1(config-if)#**lacp active**

- Use the **no** version to restore the default behavior.

lacp port-priority

- Use to set the priority for an Ethernet link in a LAG bundle.
- The member with the lowest value has the highest priority, and is selected to join the LAG bundle first.
- Valid values are in the range 0–65535.
- Example

  host1(config-if)#**lacp port-priority 100**

- Use the **no** version to restore the default value of 32768.

member-interface

- Use to add a Fast Ethernet interface or Gigabit Ethernet interface, also known as a bundle member, to a LAG bundle.
- Example

  host1(config-if)#**member-interface fastEthernet 4/0**

- Use the **no** version to remove the specified Ethernet link from the bundle.

mpls

- Use to enable, disable, or delete MPLS on an interface. MPLS is disabled by default.
- Example

  host1(config)#**mpls**

- Use the **no** version to halt MPLS on the interface and delete the MPLS interface configuration.

pppoe subinterface lag

- Use to create a PPPoE subinterface on a LAG bundle.
- Example

  host1(config-if)#**pppoe subinterface lag boston.1**

- Use the **no** version to remove the PPPoE subinterface from the LAG bundle.

virtual-router

- From Global Configuration mode, use this command to create a virtual router or access the context of a previously created virtual router or a VRF.
- Example

  host1(config)#**virtual-router boston**

- Use the **no** version of the command only to delete the VR and return the router to the default VR.

Configuring Ethernet Link Redundancy

You can use 802.3ad Link Aggregation (LAG) to configure Ethernet link redundancy for Fast Ethernet and Gigabit Ethernet interfaces. Ethernet link redundancy enables you to protect against physical link failure and account for network topology changes that redirect network traffic to redundant ports.

The following configurations are available:

- LAG to LAG—Provides redundancy capabilities for two or more ports that are assigned to a LAG. One member link is configured as the backup interface for all other ports in the LAG bundle (1:N). Traffic is not forwarded over the backup member interface; it is disabled until it takes over for an active member interface.

- LAG to non-LAG—Provides redundancy capabilities when redundant ports are connected to a bridged network that has Rapid Spanning Tree Protocol (RSTP) controlling the topology. This configuration supports only two links in the LAG.

For information about the modules that support link aggregation, see *ERX Module Guide, Appendix A, Module Protocol Support* and *E120 and E320 Module Guide, Appendix A, IOA Protocol Support*.

Ethernet Link Redundancy Configuration Models

The link connections determine the configuration model for link redundancy. The following connection types are available:

- Single-homed—Connections are between the local Ethernet interface and a single remote device. When the peer is also configured with LAG, LACP can be used to control link access.

- Dual-homed—Connections are between two separate, uncoordinated remote devices. The remote interfaces can be on the same module or on separate hardware. If LAG is not configured on the peers, LACP cannot be used to select ports; other protocols such as RSTP can be used.

The type of hardware used for connections further characterizes the single-homed and dual-homed configuration models. The following hardware types are available:

- Homogeneous—Remote interface is on another Fast Ethernet or Gigabit Ethernet port in a back-to-back router configuration of identical hardware and JUNOSe software versions. Both interfaces support the same redundant cabling and algorithm. The interfaces can be cabled on the same ports (port 0–port 0, port 1–port 1) or cross-cabled (port 0–port 1, port 1–port 0).

- Heterogeneous—Remote interface is on a different type of hardware that might or might not support redundant cabling, or on the same type of equipment with different software versions. For example, a heterogeneous configuration can include an ES2-S1 GE-4 IOA and an ES2-S1 GE-8 IOA on the E320 router, or an E-series router operating JUNOSe software connected to another vendor's router and software.

NOTE: You cannot configure link redundancy across different types of line modules in a router. You also cannot configure link redundancy across two GE-4 IOAs on the E120 router or the E320 router.

Figure 21 illustrates the configuration models for Ethernet link redundancy.

Figure 21: Ethernet Link Redundancy Configuration Models

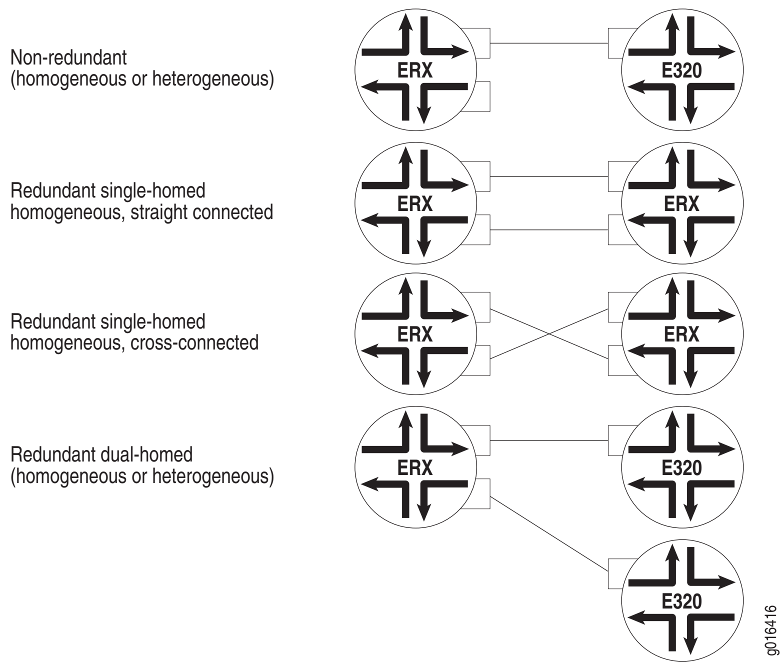

Ethernet Link Redundancy Configuration Diagrams

The diagrams in this section illustrate examples of Ethernet link redundancy configurations. The diagrams display adjacent ports bundled in a LAG.

GE-2 Line Module Configurations

These diagrams compare physical port redundancy and link redundancy on a GE-2 line module.

Figure 22 displays a GE-2 line module with physical port redundancy on both ports.

Figure 22: GE-2 Line Module Using Physical Port Redundancy

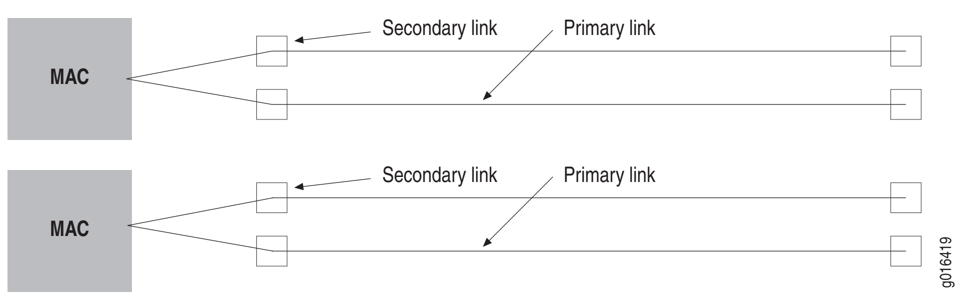

Figure 23 displays a single-homed configuration with port 0 backing up port 1 on a GE-2 line module.

Figure 23: Single-Homed GE-2 Line Module Configuration

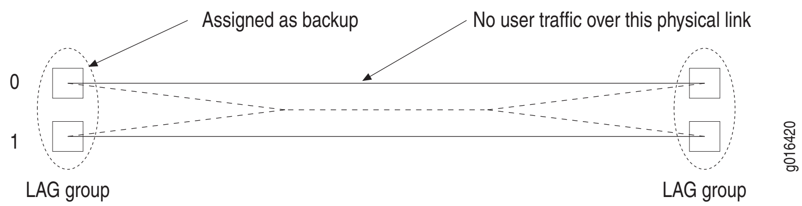

FE-8 Line Module Configurations

Figure 24 displays an FE-8 line module with a link failure in a 1:N single-homed configuration.

Figure 24: Single-Homed FE-8 Line Module Configuration (1:N)

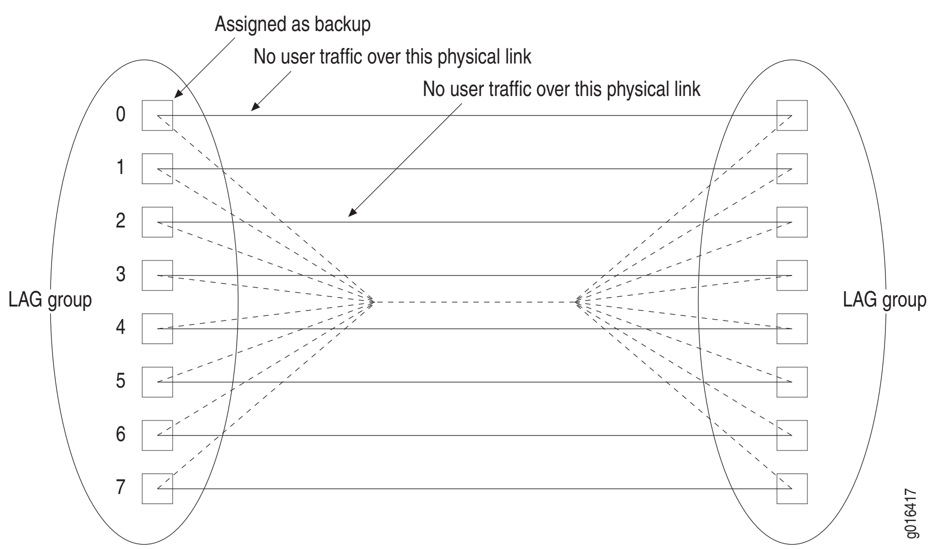

Figure 25 displays an FE-8 line module with four redundant Ethernet links in a 1:1 configuration.

Figure 25: FE-8 Line Module with 4 Redundant Ethernet Links (1:1)

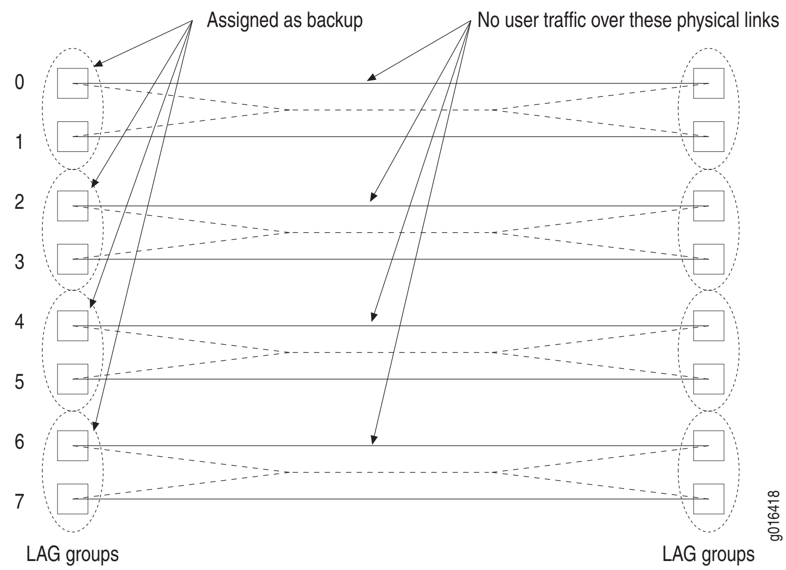

E120 and E320 Router Configurations

Figure 26 and Figure 27 display link redundancy configurations on the E120 and E320 routers.

Figure 26 displays a single-homed 1:4 configuration on an E120 router.

Figure 26: Single-Homed GE-4 IOA Configuration (1:4)

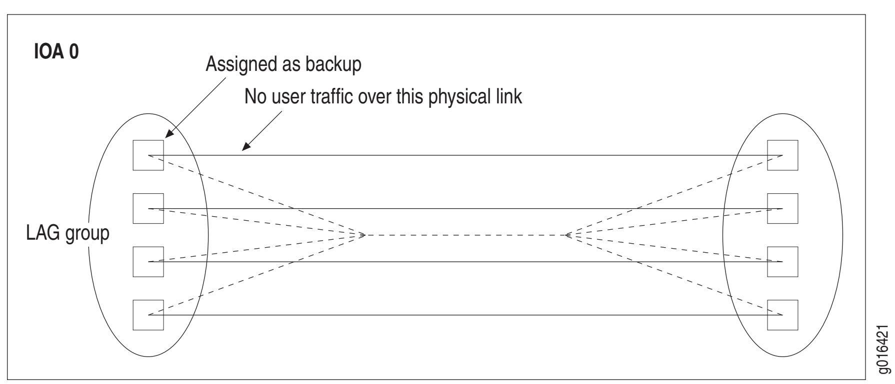

Figure 27 displays an E320 router with 1:N configuration across IOAs.

Figure 27: GE-8 IOA Configuration Across IOAs (1:N)

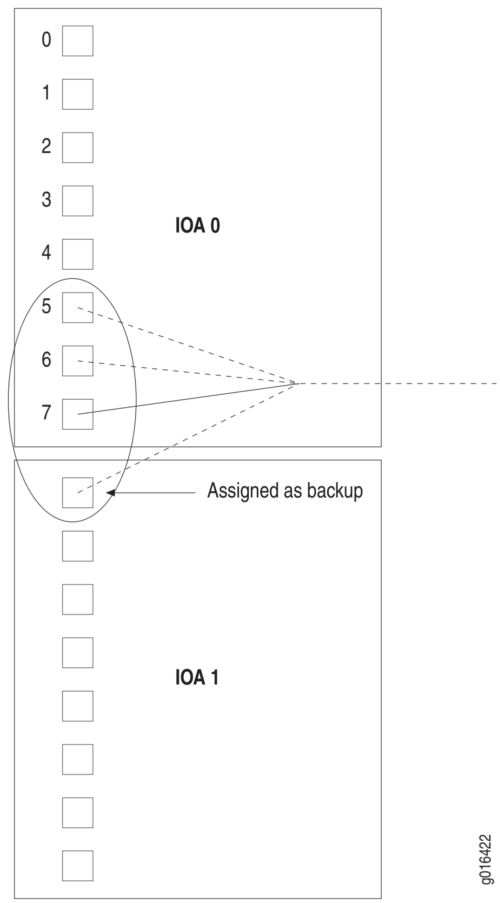

Dual-Homed Configurations with LAG Disabled

Figure 28 displays how you can configure Ethernet link redundancy with LACP disabled locally using a dual-homed configuration. LACP is disabled because there is no LAG at the peer.

Figure 28: Dual-Homed Configuration (1:1)

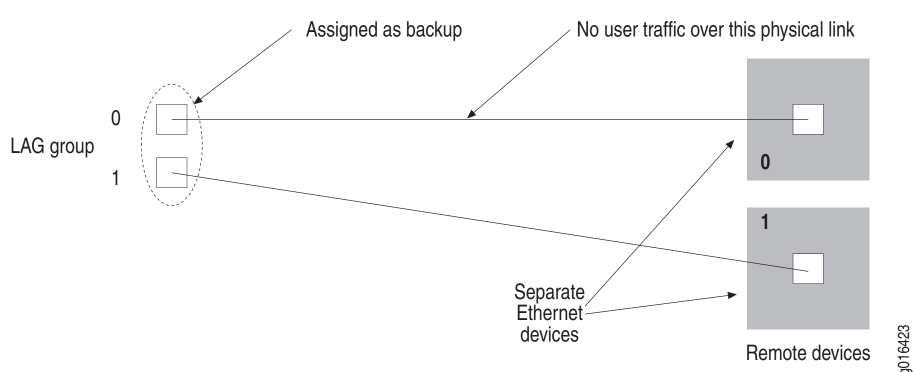

Ethernet Link Redundancy Behavior

When you create a LAG bundle, you can configure LACP with the Disabled, Passive, or Active states. For more information about these states, see *LACP* on page 206.

The following sections describe link redundancy behavior when the:

- Configuration and status of LACP changes during link failure and acquisition.
- Configuration of the endpoints of the member links is different.
- Configuration is LAG to non-LAG in an RSTP network.

Link Failure and Acquisition

Link failure on the local system occurs when the active link is no longer active. Failures can be characterized as physical link failure or virtual link failure.

Each type of link failure has different requirements for detection, failover, and link acquisition. In all cases, you configure the link to fail over when it fails by issuing the **redundant-port** command. Optionally, you can force the failover automatically by issuing the **redundant-port force-failover** command.

Protecting Against Physical Link Failure

Physical link failures can occur when a cable is cut.

To protect against physical link failure, issue the **transmitter** keyword with the **redundant-port** command to enable or disable the local redundant link. When the redundant link needs to be down, the link behavior in failure detection and failover follows a similar port redundancy scheme available with line modules such as the GE-2 line module.

When the transmitter on the remote end is enabled on the redundant link before it fails over, the local system considers the redundant link to be viable and enables the transmitter if it is disabled. If the remote end is disabled, the local end must enable the transmitter and wait for the remote end to enable.

Protecting Against Virtual Link Failure

A virtual link failure can occur when the active link is no longer used by the network because of topology changes caused by physical failure in the network. Topology changes can occur when, for example, a link is blocked because of network protocols such as RSTP blocking the port leading to selection of the redundant port connected to the receiver.

To protect against virtual link failure in conjunction with network protocols, use the **packet-sampling** keyword with the **redundant-port** command to detect link the viability. For example, when there is a network protocol decision that changes the topology and blocks a link to compensate for failures in the network, the system monitors the traffic to detect the change in network topology and fails over to the redundant port if necessary. It also determines whether the failover is successful. For more information, see *Member Link with Non-LAG Partner* on page 222.

Reverting After a Failover

When you specify the **auto-revert** keyword with the **redundant-port** command, the redundant link reverts back to redundant mode when the failed link becomes active again.

The system uses the following process when you issue the **auto-revert on** and **auto-revert off** keywords:

auto-revert on

1. An active link fails and a redundant link becomes active.
2. The original active link becomes active.
3. The original redundant link fails over to the original active link.
4. The redundant link can fail over to any other active link again.

auto-revert off

1. An active link fails and a redundant link becomes active.
2. The original active link becomes active.
3. The original redundant link remains the active link.
4. You can force the link to fail over by issuing the **redundant-port force-failover** command.

LACP Configuration and Member Link Behavior

By default, when a redundant member link is configured, the system disables LACP and the transmitter on that link.

When a member link is administratively down, the link state is operationally down at the local and remote ends, which means it does not transmit or receive PDUs.

The active link does not fail over when:

- An active link goes down and you set the redundant member link to administratively down.
- An active link is set to administratively down.

LACP configurations affect member link behavior based on the local or remote endpoint. For a remote end to include a member link in link aggregation, the two member links that are connected must have LACP configured.

Table 23 lists the acceptable configurations that enable redundant behavior for LACP modes at local and remote endpoints.

Table 23: Behavior of Member Links Using Local and Remote LACP Modes

		Remote LACP Mode		
		Disabled	**Passive**	**Active**
Local LACP Mode	Disabled	a	a	–
	Passive	a	a	a
	Active	–	a	a

Member Link with Non-LAG Partner

When a member link has a non-LAG partner, there are two separate links in a 1:1 configuration. To successfully configure this, you must disable LACP.

When a failover occurs and LACP is active, the partner might receive a new LAG ID and the LACP PDUs receive a new MAC address; therefore, the member links are not aggregated or the bundle is disabled, terminating the sessions above it.

The partner that is connected to the redundant link must not be forwarding network traffic; that is, it is either blocked through a protocol such as RSTP, or MAC address learning has selected the active port. The redundant link must not transmit over the redundant link to that MAC. The behavior of the redundant link depends on the failure detection method that is controlled by the network protocol that is blocking the port.

Ethernet Link Redundancy and RSTP

In a LAG to non-LAG configuration, you can configure redundancy capabilities when redundant ports are connected to a bridged network that has RSTP controlling the topology.

On external devices, we recommend that you configure RSTP-enabled bridged ports that are connected to the LAG interfaces as edge ports to enable the ports to transition quickly to forwarding state upon reconfiguration, and to avoid the listening and learning states required by the spanning tree protocol. The edge port designation instructs the local bridge that bridge loops do not exist through the interface, enabling it to skip the listening and learning states.

Figure 29: Dual-Homed Heterogeneous Configuration in an RSTP Network

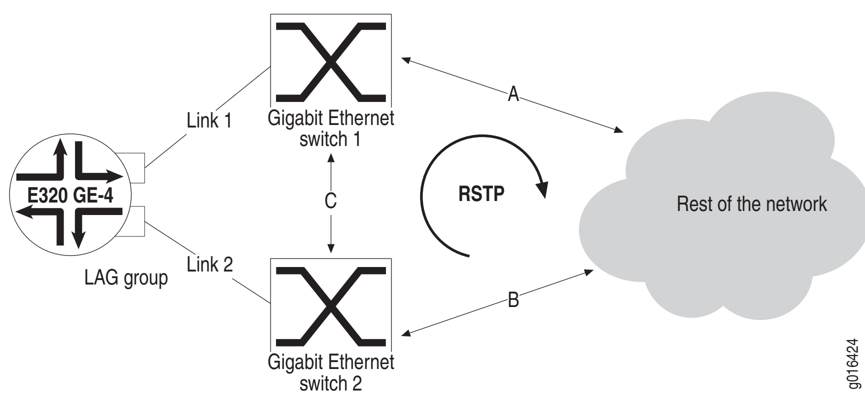

Figure 29 displays a network with RSTP enabled on Gigabit Ethernet switches 1 and 2. The local port receives bridge PDUs (BPDU), Ethernet broadcasts, and flooded unicast packets. If Link 1 is initially active and Link 2 is the backup, initial traffic destined for the LAG can be Ethernet broadcasts, PPPoE PDUs, or flooded Ethernet unicasts. The responses are only sent on the active link; in this case, Link 1.

The Ethernet network topology that is managed by RSTP learns that the MAC for the LAG group is through Link 1. Broadcasts and flooded packets are still sent on Link 2. If Link 1 is no longer viable, but has not suffered a physical failure, then that address ages out of the bridge databases and any packets directed to the LAG are flooded. The LAG detects traffic on Link 2 after the minimum delay time and then fails over.

Acquiring Initial Links

In an RSTP network, the system uses the following process for acquiring new links:

1. Based on the configuration, the system selects a link as active and the other as redundant.

2. The spanning tree converges on a topology.

3. When convergence occurs and the status of the spanning tree ports change to forwarding, network traffic appears on the links.

4. The local port detects the traffic and confirms the active member as active and the other as the redundant port. Because the initial traffic is broadcast or flooded, both ports receive the packets. However, because of the timing difference, the selected active port remains active.

Detecting Failures

In an RSTP network, the system uses the following process for detecting when the link has switched over due to topology changes:

1. BPDUs are ignored on the redundant port and system time is not retrieved. Because MAC learning forces non-flooded unicast packets to the active link, traffic to the redundant link does not receive non-flooded packets. The most recent system time is always retrieved when a network packet is received.

2. When the network cannot reach the active link because of topology changes, traffic appears on the redundant link. The redundant port detects the traffic and captures the latest timestamp. When the difference between the timestamp of the first non-bridged PDU and the time the last packet that was received on the active port is sufficiently large to account for the minimum spanning tree convergence time and latency for flooded and broadcast packets, then the port fails over.

Failing Over

In an RSTP network, the system uses the following process to fail over:

1. When the link has failed over, the system monitors the previously active port.

2. When a network packet is received on the redundant port, the system retrieves the timestamp. If the difference in timestamps between that one and the most recent on the current active port is more than the configured failover delay time, then the link fails over. If the difference is less than the delay time, the system ignores it but counts the event. If many of these transitions occur in a time period, then the system administratively brings the ports down. If no network traffic is received on either port, then failover does not occur.

Configuring Ethernet Link Redundancy

To configure Ethernet link redundancy:

1. Specify the Fast Ethernet or Gigabit Ethernet interface on which to configure a redundant link.

   host1(config)#**interface gigabitEthernet 1/1**

2. For LAG to non-LAG configurations only, specify that LACP is disabled on the port.

   host1(config-if)#**no lacp**

3. Configure a backup interface and disable LACP on it.

   host1(config)#**interface gigabitEthernet 1/0**
   host1(config-if)#**no lacp**

4. Configure a LAG interface and assign a member link to the backup interface.

   host1(config)#**interface lag myBundle**
   host1(config-if)#**member-interface gigabitEthernet 1/0**

5. Do one of the following:

   - Configure link redundancy on the port you specified in step 1.

     host1(config-if)#**redundant-port gigabitEthernet 1/1**

   - Force the port you specified in step 1 to fail over.

     host1(config-if)#**redundant-port gigabitEthernet 1/1 force-failover**

redundant-port

- Use to specify a member link in a LAG bundle as redundant.
- Use the **failover timeout** keyword to configure the amount of time between the current link event leading to failover or reversion and the previous link failover or reversion.
- Use the **packet-sampling** keyword to configure redundancy on a LAG to non-LAG application where packet sampling is used for failover detection. Use the optional **delay** keyword to control the minimum time difference to force packets on the active and redundant port to fail over.
- Use the **transmitter** keyword to enable or disable the transmitter when in redundant mode.
- Use the **auto-revert** keyword to instruct the redundant link to revert back to redundant mode when the failed link becomes active again.
- Example 1—Specifies that the Gigabit Ethernet interface in slot 4, port 0 is a redundant member interface

  host1(config-if)#**redundant-port gigabitEthernet 4/0**

- Example 2—Specifies that the Gigabit Ethernet interface in slot 1, port 1 is a redundant member interface with a packet sampling delay of 500 ms

  host1(config-if)#**redundant-port gigabitEthernet 1/1 packet-sampling delay 500**

- Use the **no** version to disable the redundant status of the member interface or disable the specified redundancy setting for the member.

redundant-port force-failover

- Use to force the specified member interface to fail over when more than one active member exists.
- Example

  host1(config)#**redundant-port gigabitEthernet 4/0 force-failover**

- There is no **no** version.

Configuring Higher-Level Protocols over Ethernet

You can configure one or more protocols over Ethernet with or without VLANs. This section focuses on non-VLAN configurations only. You can configure the following higher-level protocols on Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces:

- IP
- Point-to-Point Protocol over Ethernet (PPPoE)
- Multiprotocol Label Switching (MPLS)

The Ethernet configuration examples in this section use combinations of these protocols. Figure 30 on page 225 illustrates how different protocols can be multiplexed over a single physical link without the use of VLANs.

Figure 30: Multiplexing Multiple Protocols over a Single Physical Link

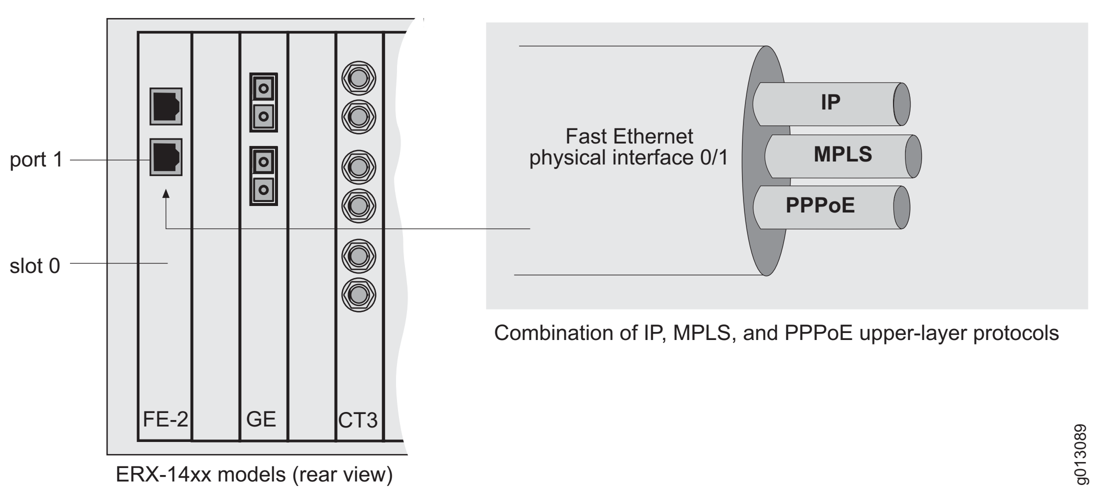

The following sections describe how to create the following common non-VLAN configurations, which you can configure on Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces:

- IP over Ethernet
- PPPoE over Ethernet
- IP over Ethernet and MPLS over Ethernet
- IP over Ethernet, MPLS over Ethernet, and PPPoE over Ethernet

Configuring IP over Ethernet

To configure IP over an Ethernet interface:

1. Specify a Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet port.

   host1(config)#**interface fastEthernet 4/1**

2. Create an IP interface.

   host1(config-if)#**ip address 192.5.127.8 255.255.255.0**

Figure 31 illustrates this configuration.

Figure 31: Example of IP over Ethernet Stacking Configuration Steps

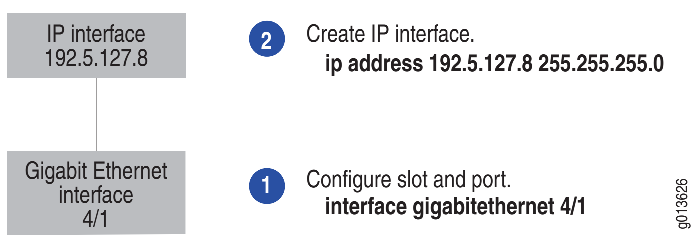

Configuring PPPoE over Ethernet

To configure PPPoE over an Ethernet interface:

1. Specify a Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet port.

   host1(config)#**interface fastEthernet 4/1**

2. Specify PPPoE as the encapsulation method on the interface.

   host1(config-if)#**pppoe**

3. Create a PPPoE subinterface.

   host1(config-if)#**pppoe subinterface fastEthernet 4/1.1**

4. Specify PPP as the encapsulation method on the interface.

   host1(config-if)#**encapsulation ppp**

5. Assign an IP address and mask.

   host1(config-if)#**ip address 164.10.6.51 255.255.255.0**

6. (Optional) Configure additional PPPoE subinterfaces by completing Steps 3 through 5 using unique numbering.

Figure 32 illustrates this configuration.

Figure 32: Example of PPPoE Stacking Configuration Steps

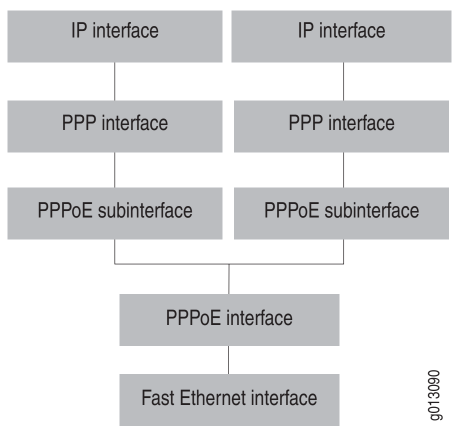

Configuring IP and MPLS over Ethernet

To configure both IP and MPLS over an Ethernet interface:

1. Specify a Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet port.

   host1(config)#**interface fastEthernet 4/0**

2. Create an IP interface.

   host1(config-if)#**ip address 192.5.127.8 255.255.255.0**

3. Create an MPLS interface.

   host1(config-if)#**mpls**

Figure 33 illustrates this configuration.

Figure 33: Example of IP and MPLS Stacking Configuration Steps

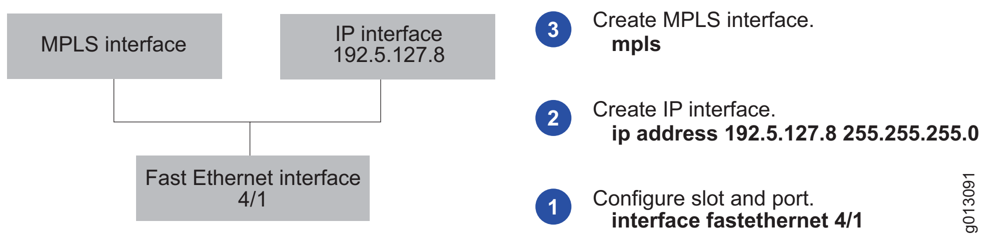

Configuring IP, MPLS, and PPPoE over Ethernet

To configure IP, MPLS, and PPPoE over an Ethernet interface:

1. Specify a Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet port.

   host1(config)#**interface fastEthernet 4/0**

2. Create an IP interface.

   host1(config-if)#**ip address 192.5.127.8 255.255.255.0**

3. Create an MPLS interface.

   host1(config-if)#**mpls**

4. Create a PPPoE interface by specifying PPPoE as the encapsulation method on the interface.

   host1(config-if)#**pppoe**

5. Create a PPPoE subinterface.

   host1(config-if)#**pppoe subinterface fastEthernet 4/1.1**

6. Specify PPP as the encapsulation method on the interface.

   host1(config-if)#**encapsulation ppp**

7. Assign an IP address and mask.

   host1(config-if)#**ip address 192.6.129.5 255.255.255.0**

8. (Optional) Configure additional PPPoE subinterfaces by completing Steps 5 through 7 using unique numbering.

Figure 34 illustrates this configuration.

Figure 34: Example of IP, MPLS, and PPPoE Stacking Configuration Steps

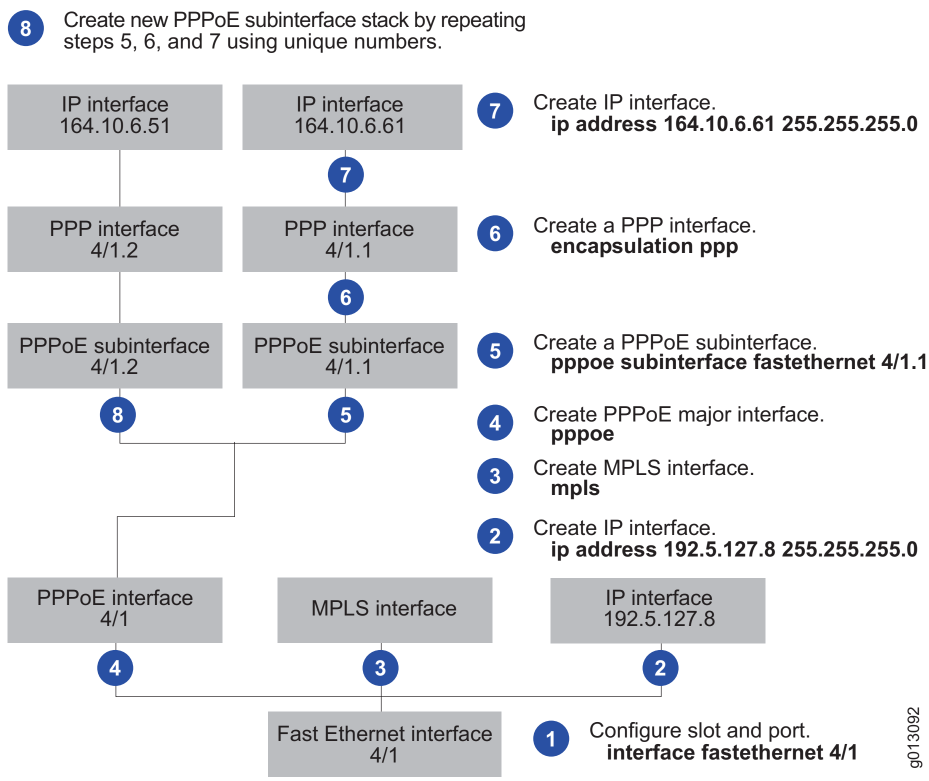

mpls

- Use to enable, disable, or delete MPLS on an interface. MPLS is disabled by default.
- Example

  host1(config)#**mpls**

- Use the **no** version to halt MPLS on the interface and delete the MPLS interface configuration.

Ethernet Link Aggregation and MPLS

CE-side load balancing in a Martini layer 2 transport environment enables an E-series router to interoperate with an 802.3ad switch in a topology designed for Ethernet link aggregation. See *JUNOSe BGP and MPLS Configuration Guide, Chapter 5, Configuring Layer 2 Services over MPLS* for more information.

Disabling Ethernet Interfaces

Use the **shutdown** command to disable an Ethernet interface.

shutdown

- Use to disable an Ethernet interface.
- Example

  host1(config-if)#**shutdown**

- Use the **no** version to restart a disabled Ethernet interface.

Monitoring Ethernet Interfaces

This section explains how to set a statistics baseline, display bit rate and packet rate statistics for VLAN subinterfaces, and use the **show** commands to display the physical characteristics and the configured settings for Ethernet interfaces.

NOTE: The E120 router and E320 router output for **monitor** and **show** commands is identical to output from other E-series routers, except that the E120 and E320 router output also includes information about the adapter identifier in the interface specifier (*slot/adapter/port*).

Setting Statistics Baselines

The router stores statistics in counters that reset only when you reboot. However, you can establish a baseline for Ethernet statistics by setting a group of reference counters to zero.

baseline interface fastEthernet | gigabitEthernet | tenGigabitEthernet

- Use to establish a baseline for Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet statistics on a line module or an SRP module.
- Use the **delta** keyword with the **show interfaces fastEthernet**, the **show interfaces gigabitEthernet**, or the **show interfaces tenGigabitEthernet** command to display baselined statistics.

Displaying Interface Rate Statistics for VLAN Subinterfaces

You can use the **monitor vlan interface** command to display bit rate and packet rate statistics over a specified time interval for one or more VLAN subinterfaces configured on the router.

To display interface rate statistics for VLAN subinterfaces:

1. Log in to the router by using a local console session or a virtual terminal (vty) session (such as a Telnet session).

   While you are using the **monitor vlan interface** command, you must keep the console or terminal session open and you cannot issue any other commands at the session during this time.

   For information about logging in to the router, see *Accessing the CLI* in *JUNOSe System Basics Configuration Guide, Chapter 2, Command-Line Interface*.

2. Access User Exec mode or Privileged Exec mode.

   For information, see *Accessing Command Modes* in *JUNOSe System Basics Configuration Guide, Chapter 2, Command-Line Interface*.

3. Specify the interface identifier for each VLAN subinterface that you want to monitor.

   host1#**monitor vlan interface fastEthernet 0/0.1 fastEthernet 4/0.1 display-time-of-day**

   For information about specifying interface identifiers for VLAN subinterfaces configured over Ethernet interfaces, see *Configuring VLANs* on page 187. For information about specifying interface identifiers for VLAN subinterfaces configured over LAG bundles, see *Configuring a VLAN Subinterface for a LAG Bundle* on page 208.

   By default, the router uses a 5-second time interval between polls to calculate bit rates and packet rates for each specified VLAN subinterface. Optionally, you can use the **load-interval** keyword to specify a nondefault time interval in the range 5–30 seconds.

   You can also include the optional **display-time-of-day** keyword to show the time of day at which the router gathers statistics for each interval. Displaying the time of day enables you to monitor when a particular VLAN subinterface is underutilized or overutilized.

4. Review the command output.

```
host1#monitor vlan interface fastEthernet 0/0.1 fastEthernet 4/0.1
display-time-of-day
                        Seconds
                        between                                          Time
       Interface         polls    Input bps/pps   Output bps/pps   (UTC)
----------------------- -------- --------------- --------------- --------
FastEthernet 0/0.1             0           --/--           --/-- 10:50:07
FastEthernet 4/0.1             0           --/--           --/-- 10:50:07
FastEthernet 0/0.1             5      120240/100      120240/100 10:50:12
FastEthernet 4/0.1             5      120000/100      120000/100 10:50:12
FastEthernet 0/0.1             5      120240/100      120240/100 10:50:17
FastEthernet 4/0.1             5      120000/100      120000/100 10:50:17
```

The router polls each VLAN subinterface at the specified load interval (the default 5-second interval in this example) to calculate and display bit rate and packet rate statistics. The first line of output for each interface always displays 0 (zero) for the number of seconds between polls, and dashes (--/--) in the Input bps/pps and Output bps/pps columns. These values indicate that the router initially takes a baseline for each interface against which to measure subsequent statistics. The router continues to display subsequent lines of output for each interface at the specified load interval until you press Ctrl + c to stop the command.

For a description of each field in the **monitor vlan interface** command output, see **monitor vlan interface** on page 232.

5. When you are finished, press Ctrl + c to stop the **monitor vlan interface** command.

```
host1#^C
```

monitor vlan interface

- Use to display bit rate and packet rate statistics over a specified time interval for one or more VLAN subinterfaces.
- You must use the **monitor vlan interface** command in a dedicated console or terminal session for the duration of the monitoring session.
- Specify the interface identifier for each VLAN subinterface that you want to monitor.
- To specify a nondefault time interval in the range 5–30 seconds at which the router calculates bit rate and packet rate statistics, use the optional **load-interval** keyword. The default time interval is 5 seconds.
- To display the time at which the router calculates bit rate and packet rate statistics for the current interval, use the optional **display-time-of-day** keyword.
- To stop the **monitor vlan interface** command, press Ctrl + c.
- Field descriptions
  - Interface—Interface identifier for the Ethernet or LAG interface on which the VLAN subinterface resides
  - Seconds between polls—Number of seconds at which the router calculates bit rate and packet rate statistics
  - Input bps/pps—Number of bits per second (bps) and packets per second (pps) received on this interface during the specified load interval
  - Output bps/pps—Number of bits per second (bps) and packets per second (pps) transmitted on this interface during the specified load interval
  - Time—Time of day, in hh:mm:ss format, at which the router calculates the bit rate and packet rate statistics for the current interval

- Example 1—Displays bit rate and packet rate statistics over the default (5-second) load interval for a single VLAN subinterface

```
host1#monitor vlan interface fastEthernet 0/0.1
                        Seconds
                        between
       Interface         polls    Input bps/pps   Output bps/pps
----------------------- -------- --------------- ---------------
FastEthernet 0/0.1             0           --/--           --/--
FastEthernet 0/0.1             5      120240/100      120240/100
FastEthernet 0/0.1             5      120000/100      120000/100
FastEthernet 0/0.1             5        92400/77        92400/77
FastEthernet 0/0.1             5        88800/74        88800/74
FastEthernet 0/0.1             5      120000/100      120000/100
host1#^C
```

- Example 2—Displays bit rate and packet rate statistics over a 10-second load interval for two VLAN subinterfaces, with the time of day that the statistics were calculated

```
host1#monitor vlan interface fastEthernet 0/0.1 fastEthernet 4/0.1
load-interval 10 display-time-of-day
                        Seconds
                        between                                     Time
       Interface         polls    Input bps/pps   Output bps/pps   (UTC)
----------------------- -------- --------------- --------------- --------
FastEthernet 0/0.1             0           --/--           --/-- 10:50:33
FastEthernet 4/0.1             0           --/--           --/-- 10:50:33
FastEthernet 0/0.1            10      120120/100      120120/100 10:50:43
FastEthernet 4/0.1            10      120000/100      120000/100 10:50:43
FastEthernet 0/0.1            10      120000/100      120000/100 10:50:53
FastEthernet 4/0.1            10      120000/100      120000/100 10:50:53
host1#^C
```

- There is no **no** version.

Using Ethernet show Commands

Use the **show** commands described in this section to display information about your Ethernet configuration and to monitor Ethernet interfaces.

show interfaces fastEthernet

- Use to display the status of Fast Ethernet interfaces, VLAN subinterfaces, or S-VLAN subinterfaces.
- You can specify the following keywords:
  - **delta**—Specifies that baselined statistics are to be shown
  - **brief**—Displays the operational status of all configured interfaces
- Field descriptions
  - FastEthernet *interfaceSpecifier*—Status of the hardware on this interface
    - up—Hardware is operational
    - down—Hardware is not operational

- Administrative status—Operational state that you configured for this interface
  - up—Interface is enabled
  - down—Interface is disabled
- Hardware—Type of MAC device on this interface
- Address—MAC address of the processor on this interface
- MAU—Type of medium attachment unit (MAU) on the physical port:
  - 10BASE-T (10 Mbps)
  - 100BASE-TX (100 Mbps)
  - 100BASE-FX-MM (100 Mbps) with the distance appearing after the type
  - 100BASE-LX-SM (100 Mbps)
  - SFP (Empty)—SFPs that are empty
  - SFP (Non-compliant Juniper Part)—SFPs that are installed in the FE-8 I/O module and do not have a Juniper Networks part number programmed
- MTU—Size of the MTU for this interface
  - Operational—Size of the largest packet processed
  - Administrative—Setting for MTU size that you specified
- Duplex Mode—Duplex option for this interface
  - Operational—Duplex option currently used
  - Administrative—Setting for duplex that you specified
- Speed—Line speed for this interface
  - Operational—Current rate at which packets are processed
  - Administrative—Setting for line speed
  - 5 minute input rate—Data rates based on traffic received in the last 5 minutes
  - 5 minute output rate—Data rates based on traffic sent in the last 5 minutes
- In—Analysis of inbound traffic on this interface
  - Bytes—Number of bytes received in error-free packets
  - Unicast—Number of unicast packets received
  - Multicast—Number of multicast packets received
  - Broadcast—Number of broadcast packets received
  - Errors—Total number of errors in all received packets; some packets might contain more than one error
  - Discards—Total number of discarded incoming packets
  - Mac Errors—Number of incoming packets discarded because of MAC sublayer failures
  - Alignment—Number of incomplete octets received

  - CRC—Number of packets discarded because the checksum the router computed from the data does not match the checksum generated by the originating devices
  - Too Longs—Number of packets discarded because the size exceeded the MTU
  - Symbol Errors—Number of symbols received that the router did not correctly decode
- Out—Analysis of outbound traffic on this interface
  - Bytes—Number of bytes sent
  - Unicast—Number of unicast packets sent
  - Multicast—Number of multicast packets sent
  - Broadcast—Number of broadcast packets sent
  - Errors—Total number of errors in all transmitted packets; some packets might contain more than one error
  - Discards—Total number of discarded outgoing packets
  - Mac Errors—Number of outgoing packets discarded because of MAC sublayer failures
  - Deferred—Number of packets that the router delayed sending because the line was busy. In half duplex mode, a high number of deferrals means the link is very busy with traffic from other stations. In full duplex mode, when the link is always available for transmission, this number is zero.
  - No Carrier—Number of packets sent when carrier sense was unavailable
- Collisions—Analysis of the collisions that occurred
  - Single—Number of packets sent after one collision
  - Multiple—Number of packets sent after multiple collisions
  - Late—Number of packets aborted during sending because of collisions after 64 bytes
  - Excessive—Number of packets not sent because of too many collisions
- ARP Statistics—Analysis of ARP traffic on this interface; In fields are for traffic received on the interface and Out fields are for traffic sent on the interface
  - ARP requests—Number of ARP requests
  - ARP responses—Number of ARP responses
  - Errors—Total number of errors in all ARP packets
  - Discards—Total number of discarded ARP packets
- queue—Hardware packet queue associated with the specified traffic class and interface
  - Queue length—Length of the queue, in bytes
  - Forwarded packets, bytes—Number of packets and bytes that were forwarded on this queue

        - Dropped committed packets, bytes—Number of committed packets and bytes that were dropped
        - Dropped conformed packets, bytes—Number of conformed packets and bytes that were dropped
        - Dropped exceeded packets, bytes—Number of exceeded packets and bytes that were dropped

- Field descriptions when you display the status of a Fast Ethernet VLAN or S-VLAN subinterface
    - *Subinterface number*—Location of the subinterface that carries the VLAN or S-VLAN traffic
    - Administrative status—Operational state that you configured for this interface; up or down
    - VLAN ID—Domain number of the VLAN
    - SVLAN ID—Domain number of the stacked VLAN
    - Ethertype—Ethertype assignment for the S-VLAN subinterface, 0x8100, 0x88a8, or 0x9100; 0x9100 is the default
    - In—Analysis of inbound traffic on this interface
        - Bytes—Number of bytes received on the VLAN or S-VLAN subinterface
        - Packets—Sum of all unicast, broadcast, and multicast packets received on the VLAN or S-VLAN subinterface
        - Multicast—Number of multicast packets received on the VLAN or S-VLAN subinterface
        - Broadcast—Number of broadcast packets received on the VLAN or S-VLAN subinterface
        - Errors—Total number of errors in all received packets; some packets might contain more than one error
        - Discards—Total number of discarded incoming packets
    - Out—Analysis of outbound traffic on this interface
        - Bytes—Number of bytes sent on the VLAN or S-VLAN subinterface
        - Packets—Number of packets sent on the VLAN or S-VLAN subinterface
        - Multicast—Number of multicast packets sent on the VLAN or S-VLAN subinterface
        - Broadcast—Number of broadcast packets sent on the VLAN or S-VLAN subinterface
        - Errors—Total number of errors in all transmitted packets; note that some packets might contain more than one error
        - Discards—Total number of discarded outgoing packets

- Example 1—Displays the status of a Fast Ethernet interface

```
host1:vr2#show interfaces fastEthernet 2/0
FastEthernet2/0 is Up, Administrative status is Up
  Hardware is Intel 21440, address is 0090.1a10.0552
  MAU is 10BASE-T
  MTU: Operational 1518, Administrative 1518
  Duplex Mode: Operational Full Duplex, Administrative Auto Negotiate
  Speed: Operational 100 Mbps, Administrative Auto Negotiate

  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec

  In: Bytes 39256, Unicast 612
   Multicast 0, Broadcast 0
   Errors 0, Discards 0, Mac Errors 0, Alignment 0
   CRC 0, Too Longs 0, Symbol Errors 0
  Out: Bytes 4579036, Unicast 610
   Multicast 0, Broadcast 70932
   Errors 0, Discards 0, Mac Errors 0, Deferred 0, No Carrier 3
   Collisions: Single 0, Multiple 0, Late 0, Excessive 0
 ARP Statistics:
  In: ARP requests 0, ARP responses 0
   Errors 0, Discards 0
  Out: ARP requests 0, ARP responses 0
   Errors 0, Discards 0
Administrative qos-shaping-mode: none
Operational qos-shaping-mode: none

queue 0: traffic class control, bound to FastEthernet2/0
  Queue length 0 bytes
  Forwarded packets 1, bytes 46
  Dropped committed packets 0, bytes 0
  Dropped conformed packets 0, bytes 0
  Dropped exceeded packets 0, bytes 0
```

- Example 2—Displays the status of a Fast Ethernet VLAN subinterface

```
host1:vr2#show interfaces fastEthernet 8/3.1
FastEthernet8/3.1 is Up, Administrative status is Up
  VLAN ID: 10, address 0090.5e00.0001

In: Bytes 39256, Packets 612
   Multicast 0, Broadcast 0
   Errors 0, Discards 0
Out: Bytes 4536220, Packets 70873
   Multicast 0, Broadcast 70258
   Errors 0, Discards 0
ARP Statistics:
  In: ARP requests 1, ARP responses 0
   Errors 0, Discards 0
  Out: ARP requests 1, ARP responses 0
   Errors 0, Discards 0
```

- Example 3—Displays the status of a Fast Ethernet S-VLAN subinterface

```
host1:vr2#show interfaces fastEthernet 0/0.1
FastEthernet0/0.1 is Up, Administrative status is Up
  SVLAN ID: 1, VLAN ID: 0, Ethertype 0x9100

  In: Bytes 39256, Packets 612
   Multicast 0, Broadcast 0
   Errors 0, Discards 0
  Out: Bytes 4536220, Packets 70873
   Multicast 0, Broadcast 70258
   Errors 0, Discards 0
ARP Statistics:
  In: ARP requests 0, ARP responses 0
   Errors 0, Discards 0
  Out: ARP requests 0, ARP responses 0
   Errors 0, Discards 0
```

show interfaces gigabitEthernet
show interfaces tenGigabitEthernet

- Use to display the status of Gigabit Ethernet interfaces, 10-Gigabit Ethernet interfaces, VLAN subinterfaces, or S-VLAN subinterfaces.
- You can specify the following keywords:
  - **delta**—Specifies that baselined statistics are to be shown
  - **brief**—Displays the operational status of all configured interfaces
- Field descriptions
  - GigabitEthernet or tenGigabitEthernet *interfaceSpecifier*—Status of the hardware on this interface
    - up—Hardware is operational
    - down—Hardware is not operational
  - Administrative status—Operational state that you configured for this interface
    - up—Interface is enabled
    - down—Interface is disabled
  - Hardware—Type of MAC device on this interface
  - Address—MAC address of the processor on this interface
  - MAU—Type of medium attachment unit (MAU) on the primary and secondary physical ports:
    - SFP—1000BASE-LH, 1000BASE-SX, 1000BASE-ZX; for SFPs that are empty, SFP (Empty) appears in this field; for SFPs that are installed in the OC3-2 GE APS I/O module and do not have a Juniper Networks part number programmed, SFP (GE Compliant) appears in this field
    - XFP—10GBASE-SR (10 Gbps), 10GBASE-LR (10 Gbps), 10GBASE-ER (10 Gbps); for XFPs that are empty, XFP (Empty) appears in this field

- MTU—Size of the MTU for this interface
  - Operational—Size of the largest packet processed
  - Administrative—Setting for MTU size that you specified
- Duplex Mode—Duplex option for this interface
  - Operational—Duplex option currently used
  - Administrative—Setting for duplex that you specified
- Speed—Line speed for this interface
  - Operational—Current rate at which packets are processed
  - Administrative—Setting for line speed that you specified
- Link —Link information for this interface
  - Operational Link Selected—Port that the I/O module is currently using: primary or secondary
  - Administrative link selected—Port that the I/O module is configured to use:
    - primary—Only primary port is configured to operate
    - secondary—Only redundant port is configured to operate
    - automatically—Software controls port redundancy automatically
- Primary link selected x times—Number of times that the I/O has used the primary port since the module was last rebooted
- Secondary link selected x times—Number of times that the I/O has used the secondary port since the module was last rebooted
- Primary/Secondary link signal detected, Primary/Secondary link signal not detected—Specifies the port (primary or secondary) on which the router detects a signal
- 5 minute input rate—Data rates based on the traffic received in the last 5 minutes
- 5 minute output rate—Data rates based on the traffic sent in the last 5 minutes
- In—Analysis of inbound traffic on this interface
  - Bytes—Number of bytes received in error-free packets
  - Unicast—Number of unicast packets received
  - Multicast—Number of multicast packets received
  - Broadcast—Number of broadcast packets received
  - Errors—Total number of errors in all received packets; some packets might contain more than one error
  - Discards—Total number of discarded incoming packets
  - Mac Errors—Number of incoming packets discarded because of MAC sublayer failures
  - Alignment—Number of incomplete octets received

  - CRC—Number of packets discarded because the checksum that the router computed from the data does not match the checksum generated by the originating devices
  - Too Longs—Number of packets discarded because the size exceeded the MTU
  - Symbol Errors—Number of symbols received that the router did not correctly decode
- Out—Analysis of outbound traffic on this interface
  - Bytes—Number of bytes sent
  - Unicast—Number of unicast packets sent
  - Multicast—Number of multicast packets sent
  - Broadcast—Number of broadcast packets sent
  - Errors—Total number of errors in all transmitted packets; note that some packets might contain more than one error
  - Discards—Total number of discarded outgoing packets
  - Mac Errors—Number of outgoing packets discarded because of MAC sublayer failures
  - Deferred—Number of packets that the router delayed sending because the line was busy. In half duplex mode, a high number of deferrals means the link is very busy with traffic from other stations. In full duplex mode, when the link is always available for transmission, this number is zero.
  - No Carrier—Number of packets sent when carrier sense was unavailable
- Collisions—Analysis of the collisions that occurred
  - Single—Number of packets sent after one collision
  - Multiple—Number of packets sent after multiple collisions
  - Late—Number of packets aborted during sending because of collisions after 64 bytes
  - Excessive—Number of packets not sent because of too many collisions
- ARP Statistics—Analysis of ARP traffic on this interface; In fields are for traffic received on the interface and Out fields are for traffic sent on the interface
  - ARP requests—Number of ARP requests
  - ARP responses—Number of ARP responses
  - Errors—Total number of errors in all ARP packets
  - Discards—Total number of discarded ARP packets
- queue—Hardware packet queue associated with the specified traffic class and interface
  - Queue length—Length of the queue, in bytes
  - Forwarded packets, bytes—Number of packets and bytes that were forwarded on this queue

    - Dropped committed packets, bytes—Number of committed packets and bytes that were dropped
    - Dropped conformed packets, bytes—Number of conformed packets and bytes that were dropped
    - Dropped exceeded packets, bytes—Number of exceeded packets and bytes that were dropped

- Field descriptions when you display the status of a Gigabit Ethernet or 10-Gigabit Ethernet VLAN or S-VLAN subinterface
  - *Subinterface number*—Location of the subinterface that carries the VLAN or S-VLAN traffic
  - Administrative status—Operational state that you configured for this interface; up or down
  - VLAN ID—Domain number of the VLAN
  - SVLAN ID—Domain number of the stacked VLAN
  - Ethertype—Ethertype assignment for the S-VLAN subinterface, 0x8100, 0x88a8, or 0x9100; 0x9100 is the default
  - In—Analysis of inbound traffic on this interface
    - Bytes—Number of bytes received on the VLAN or S-VLAN subinterface
    - Packets—Sum of all unicast, broadcast, and multicast packets received on the VLAN or S-VLAN subinterface
    - Multicast—Number of multicast packets received on the VLAN or S-VLAN subinterface
    - Broadcast—Number of broadcast packets received on the VLAN or S-VLAN subinterface
    - Errors—Total number of errors in all received packets; some packets might contain more than one error
    - Discards—Total number of discarded incoming packets
  - Out—Analysis of outbound traffic on this interface
    - Bytes—Number of bytes sent on the VLAN or S-VLAN subinterface
    - Packets—Number of packets sent on the VLAN or S-VLAN subinterface
    - Multicast—Number of multicast packets sent on the VLAN or S-VLAN subinterface
    - Broadcast—Number of broadcast packets sent on the VLAN or S-VLAN subinterface
    - Errors—Total number of errors in all transmitted packets; some packets might contain more than one error
    - Discards—Total number of discarded outgoing packets
- Example 1—Displays the status of a Gigabit Ethernet interface

```
host1:vr2#show interfaces gigabitEthernet 10/2
ERX-40-20-43#show int gigabitEthernet 10/2
GigabitEthernet10/2 is Down, Administrative status is Up
  Hardware is SEEQ 8101, address is 0090.1a01.0cc8
  Primary MAU is 1000BASE-SX, secondary MAU is SFP (Empty)
  MTU: Operational 1518, Administrative 1518
```

```
  Duplex Mode: Operational Full Duplex, Administrative Auto Negotiate
  Speed: Operational 1000 Mbps, Administrative Auto Negotiate
  Link: Operational Secondary Link Selected,
        Administrative Link Selected Automatically
  Link Failover Timeout: Operational 652 ms, Administrative default
  Primary link selected 6302 times, Secondary link selected 6302 times
  Primary link signal detected, Secondary link signal detected

  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec

  In: Bytes 0, Unicast 0
   Multicast 0, Broadcast 0
   Errors 0, Discards 0, Mac Errors 0, Alignment 0
   CRC 0, Too Longs 0, Symbol Errors 0
  Out: Bytes 0, Unicast 0
   Multicast 0, Broadcast 0
   Errors 0, Discards 0, Mac Errors 0, Deferred 0, No Carrier 0
   Collisions: Single 0, Multiple 0, Late 0, Excessive 0
 ARP Statistics:
  In: ARP requests 0, ARP responses 0
   Errors 0, Discards 0
  Out: ARP requests 0, ARP responses 0
   Errors 0, Discards 0
Administrative qos-shaping-mode: none
Operational qos-shaping-mode: none

queue 0: traffic class control, bound to GigabitEthernet10/2
  Queue length 0 bytes
  Forwarded packets 0, bytes 0
  Dropped committed packets 0, bytes 0
  Dropped conformed packets 0, bytes 0
  Dropped exceeded packets 0, bytes 0
```

- Example 2—Displays the status of a Gigabit Ethernet VLAN subinterface

```
host1:vr2#show interfaces gigabitEthernet 2/0.1
GigabitEthernet2/0.1 is Up, Administrative status is Up
  VLAN ID: 10, address 0090.5e00.0001

  In: Bytes 2357, Packets 23
   Multicast 0, Broadcast 0
   Errors 0, Discards 0
  Out: Bytes 4872, Packets 57
   Multicast 0, Broadcast 0
   Errors 0, Discards 0
ARP Statistics:
  In: ARP requests 0, ARP responses 0
   Errors 0, Discards 0
  Out: ARP requests 0, ARP responses 0
   Errors 0, Discards 0
```

- Example 3—Displays the status of a Gigabit Ethernet S-VLAN subinterface

```
host1:vr2#show interfaces gigabitEthernet 2/0.2
GigabitEthernet2/0.2 is Up, Administrative status is Up
  SVLAN ID: 10, VLAN ID: 100, Ethertype 0x9100

  In: Bytes 2357, Packets 23
   Multicast 0, Broadcast 0
   Errors 0, Discards 0
  Out: Bytes 4872, Packets 57
   Multicast 0, Broadcast 57
```

```
ARP Statistics:
  In: ARP requests 0, ARP responses 0
   Errors 0, Discards 0
  Out: ARP requests 0, ARP responses 0
   Errors 0, Discards 0
```

show interfaces lag

- Use to display information about a specified Ethernet member link in an IEEE 802.3ad link aggregation group (LAG) bundle.
- Specify either the Fast Ethernet or Gigabit Ethernet interface type when issuing this command:

  host1(config):**show interfaces *interfaceType interfaceSpecifier* lag**

- Field descriptions
  - *interfaceSpecifier*—Status of the hardware on this interface
    - Up—Hardware is operational
    - Down—Hardware is not operational
  - Administrative status—Operational state that you configured for this interface
  - Member—Membership status of the Ethernet link
  - LACP—Status of LACP configuration for the Ethernet link
    - active—Ethernet link participates in the protocol regardless of whether its Partner member link is set to active or passive LACP PDU participation
    - passive—Ethernet link transmits LACP PDUs only when it receives LACP PDUs from its Partner member link
  - mux state—Status of collecting and distributing at the Mux state machine
    - collecting/distributing—Ethernet link is actively collecting incoming frames and distributing outgoing frames
    - detached—Ethernet link is detached from the LAG bundle due to protocol changes or system constraints
    - waiting—Ethernet link is waiting to attach to a LAG bundle
  - LACP state
    - active—Actor link actively particulates in LACP
    - passive—Actor link transmits LACP PDUs
    - timeout—Timeout control value; this value is not configurable and is set to long timeout (30 seconds)
    - aggregatable—Actor link can be aggregated
    - individual—Actor link cannot be aggregated; must operate as an individual link
    - in-sync—Actor link has joined the correct LAG bundle
    - out-of-sync—Actor link is unable to join the correct LAG bundle

  - collecting—Actor link is actively collecting incoming frames; if this field does not appear, the Actor link is not actively collecting incoming frames
  - distributing—Actor link is actively distributing outgoing frames; if this field does not appear, the Actor link is not actively distributing outgoing frames
  - defaulted—Actor link is using defaulted operational information about the Partner link that was administratively configured for Partner; if this field does not appear, the operational information about the Partner link has been received by the Actor link in an LACP PDU
  - expired—Actor link's receive machine is expired; if this field does not appear, the Actor link's receive machine is active
- port—Port number assigned to the Ethernet link by the Actor link
- priority—Priority assigned to this Ethernet link by the Actor link
- Key—Operational key value assigned to the Ethernet link by the Actor link
- System Priority—Priority assigned to the Ethernet link by the system
- System MAC Address—MAC address assigned to the Actor link
- Partner—Status of the Partner link
  - active— Partner link participates in the LACP
  - passive—Partner link transmits LACP PDUs
  - timeout—Timeout control value; short timeout or long timeout
  - aggregatable—Partner link can be aggregated
  - individual—Partner link cannot be aggregated
  - in-sync—Partner link has joined the correct LAG bundle
  - out-of-sync—Partner link has joined the incorrect LAG bundle
  - collecting—Partner link is actively collecting incoming frames; if this field does not appear, the Partner link is not actively collecting incoming frames
  - distributing—Partner link is actively distributing outgoing frames; if this field does not appear, the Partner link is not actively distributing outgoing frames
  - defaulted—Partner link is using defaulted operational information about the Partner link that was administratively configured for Partner; if this field does not appear, the operational information about the Partner link has been received by the Actor link in an LACP PDU
  - expired—Partner link's receive machine is expired; if this field does not appear, the Partner link's receive machine is active
  - port—Port number assigned to the Ethernet link by the Partner link
  - priority—Priority assigned to the Ethernet link by the Partner link
  - key—Operational key value assigned to the Ethernet link by the Partner link
  - age—Number of seconds since last LACP was received

    - System Priority—Priority assigned to the Ethernet link by the Partner link's system
    - System MAC Address—MAC address assigned to the Partner link by the system
  - LACP packets—Number of transmitted and received LACP packets
  - Marker Protocol request packets—Number of Marker Protocol packets requested to verify transmissions
  - Marker Protocol response packets—Number of Marker Protocol response packets that verified transmissions
  - Discarded—Number of invalid LACP packets
- Example

```
host1#show interfaces fastEthernet 4/0 lag
FastEthernet4/0 is Up, Administrative status is Up
 Member of Lag boston
 LACP passive, mux state collecting/distributing
 LACP state (0x3c) passive, long timeout, aggregatable, in-sync, collecting,
distributing
 port 0 priority 32768 key 8
 System Priority 32768 System MAC Address is 0090.1a40.2043
 Partner: state (0x3d) active, short timeout, aggregatable, in-sync,
collecting, distributing
  port 0 priority 32768 key 8 age 25
  System Priority 32768 System MAC Address is 0090.1a40.2043

 LACP packets: received 8, transmitted 7
 Marker Protocol request packets: received 0, transmitted 0
 Marker Protocol response packets: received 0, transmitted 0
 Discarded 0, unknown protocol received 0
```

show interfaces lag members

- Use to display information about the Ethernet member links in all IEEE 802.3ad link aggregation group (LAG) bundles configured on the router, or about the member links in a specified IEEE 802.3ad LAG bundle.
- Field descriptions
  - Lag—Name of the LAG bundle
  - Administrative status—Operational state that you configured for the LAG
  - Member-interface—Status of the member interface in the bundle
    - *Interface Specifier*—Status of the hardware on this interface (up or down)
    - LACP active—Ethernet link participates in the protocol regardless of whether its Partner member link is set to active or passive LACP PDU participation
    - LACP passive—Ethernet link transmits LACP PDUs only when it receives LACP PDUs from its Partner link

    - collecting/distributing—Ethernet link is actively collecting incoming frames and distributing outgoing frames
    - detached—Ethernet link is detached from the LAG bundle due to protocol changes or system constraints
    - waiting—Ethernet link is waiting to attach to a LAG bundle

- Example

```
host1#show interfaces boston lag members

Lag bostonBundle is Up, Administrative status is Up
   Member-interface FastEthernet0/0 is Up
      (LACP active, state collecting/distributing)
   Member-interface FastEthernet0/5 is Up
      (LACP active, state collecting/distributing)

Lag actonBundle is Up, Administrative status is Up
   Member-interface FastEthernet4/0 is Up
      (LACP passive, state collecting/distributing)
   Member-interface FastEthernet4/6 is Up
      (LACP passive, state collecting/distributing)
2 lag interfaces found
```

show ip mac-validate interface

- Use to display the status of the MAC address validation on the physical interface.
- Field descriptions
    - FastEthernet *interfaceSpecifier*—On the ERX-14xx models, ERX-7xx models, and ERX-310 router, the Fast Ethernet or Gigabit Ethernet interface *slot/port*; on the E120 and E320 routers, the Gigabit Ethernet or 10-Gigabit Ethernet interface *slot/adapter/port*
    - Keyword assigned to interface—Options: Strict or Loose
    - Address—IP address of the entry
    - Hardware Addr—Physical (MAC) address of the entry
- Example

```
host1:boston#show ip mac-validate interface fastEthernet 11/0
FastEthernet11/0:  Strict

         Address     Hardware Addr
         3.3.3.3     0090.1a30.3365
         4.4.4.4     0090.1a30.3368
```

show vlan subinterface

- Use to display configuration and status information for a specified VLAN subinterface or for all VLAN subinterfaces configured on the router.
- Use the **summary** keyword to display only the counts of all VLAN subinterfaces and VLAN major interfaces configured on the router.
- Use the **mac-address** keyword to display information about the VLAN subinterfaces that were configured with unique MAC addresses.
- Use the **vlan** or **svlan** keywords to display information about specific S-VLAN IDs or VLAN IDs.

- Field descriptions
  - Interface—Type and specifier of the VLAN subinterface
  - Status—Status of the VLAN subinterface: up, down, dormant, lowerLayerDown, absent
  - MTU—Maximum allowable size (in bytes) of the maximum transmission unit (MTU) for the VLAN subinterface
  - Svlan Id—S-VLAN ID value, if configured
  - Vlan Id—VLAN ID value for the VLAN subinterface
  - Ethertype—S-VLAN Ethertype value, if configured
  - Type—Type of VLAN subinterface
    - Static—VLAN or S-VLAN subinterface was configured statically
    - Dynamic—VLAN or S-VLAN subinterface was configured dynamically
  - In—Analysis of inbound traffic on this interface
    - Bytes—Number of bytes received on the VLAN or S-VLAN subinterface
    - Packets—Sum of all unicast, broadcast, and multicast packets received on the VLAN or S-VLAN subinterface
    - Multicast—Number of multicast packets received on the VLAN or S-VLAN subinterface
    - Broadcast—Number of broadcast packets received on the VLAN or S-VLAN subinterface
    - Errors—Total number of errors in all received packets; some packets might contain more than one error
    - Discards—Total number of discarded incoming packets
  - Out—Analysis of outbound traffic on this interface
    - Bytes—Number of bytes sent on the VLAN or S-VLAN subinterface
    - Packets—Number of packets sent on the VLAN or S-VLAN subinterface
    - Multicast—Number of multicast packets received on the VLAN or S-VLAN subinterface
    - Broadcast—Number of broadcast packets received on the VLAN or S-VLAN subinterface
    - Errors—Total number of errors in all transmitted packets; some packets might contain more than one error
    - Discards—Total number of discarded outgoing packets

- ARP Statistics—Analysis of ARP traffic on this interface; In fields are for traffic received on the interface and Out fields are for traffic sent on the interface
  - ARP requests—Number of ARP requests
  - ARP responses—Number of ARP responses
  - Errors—Total number of errors in all ARP packets
  - Discards—Total number of discarded ARP packets
- Total VLAN interfaces—Total numbers of VLAN subinterfaces and VLAN major interfaces configured on the router; this is the only field that appears when you specify the **summary** keyword

- Example 1—Displays full status and configuration information for all VLAN subinterfaces configured on the router

```
host1#show vlan subinterface
        Interface           Status    MTU  Svlan Id Vlan Id  Ethertype  Type
------------------------- ---------- ---- -------- -------- --------- -------
ATM 3/0.1.2               Up         1522 ----     11       ----      Static
ATM 3/0.1.3               Up         1522 ----     12       ----      Static
ATM 3/1.1.1               Up         1522 ----     13       ----      Static
ATM 3/1.1.2               Up         1522 ----     14       ----      Static
ATM 3/2.1.1               Down       1526 4        255      0x9100    Static
FastEthernet 4/5.1        Up         1522 ----     1        ----      Dynamic
6 vlan subinterfaces found
```

- Example 2—Displays full status and configuration information for the specified VLAN subinterface

```
host1#show vlan subinterface fastEthernet 0/0.1
        Interface           Status    MTU  Svlan Id Vlan Id  Ethertype  Type
------------------------- ---------- ---- -------- -------- --------- -------
FastEthernet 0/0.1        Up         1526        1        0    0x9100 Static

In: Bytes 39256, Packets 612
   Multicast 0, Broadcast 0
   Errors 0, Discards 0
  Out: Bytes 4538652, Packets 70911
   Multicast 0, Broadcast 70296
   Errors 0, Discards 0
ARP Statistics:
  In: ARP requests 0, ARP responses 0
   Errors 0, Discards 0
  Out: ARP requests 0, ARP responses 0
   Errors 0, Discards 0
```

- Example 3—Displays only brief summary information for all VLAN subinterfaces configured on the router

```
host1#show vlan subinterface summary
Total VLAN interfaces: 6 subinterfaces, 3 major interfaces
```

- Example 4—Displays full status and configuration information for all VLAN subinterfaces configured with a unique MAC address

```
host1#show vlan subinterface mac-address
        Interface         Svlan Id  Vlan Id    MAC Address
------------------------  --------  --------  --------------
FastEthernet 4/0.25           ----        25  0090.dfad.2abd
FastEthernet 4/0.10050           1        50  0090.adad.0abd
2 vlan subinterfaces found
```

- Example 5—Displays full status and configuration information for a VLAN subinterface on a LAG bundle

```
host1#show vlan subinterface lag boston.1
        Interface          Status    MTU  Svlan Id Vlan Id  Ethertype  Type
------------------------ ---------- ---- -------- -------- --------- -------
lag boston.1             Up         1522 ----     1        ----      Static
```

- Example 6—Displays full status and configuration information for the specified S-VLAN ID

```
host1#show vlan subinterface svlan 100 53
        Interface          Status    MTU  Svlan Id Vlan Id  Ethertype  Type
------------------------ ---------- ---- -------- -------- --------- -------
FastEthernet 0/0.1       Up         1526 100      53       0x9100    Static
FastEthernet 4/6.1       Up         1526 100      53       0x9100    Dynamic
2 vlan subinterfaces found
```