Chapter 6
Configuring 802.3ad Link Aggregation and Link Redundancy

This chapter describes how to configure 802.3ad link aggregation and link redundancy on E-series routers.

This chapter contains the following sections:

802.3ad Link Aggregation for Ethernet Overview

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

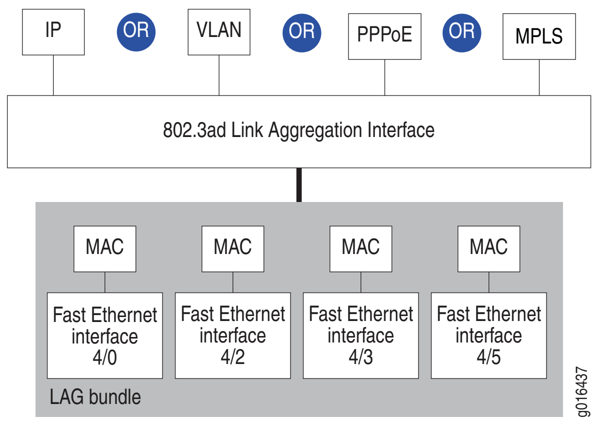

For information about configuring higher-level protocols over VLANs, see *Chapter 5, Configuring VLAN and S-VLAN Subinterfaces*.

NOTE: On the ES2 10G LM and ES2-S1 GE-8 IOA combination, you can only configure IP or VLAN over a LAG bundle.

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*.

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.

802.3ad Link Aggregation Platform Considerations

You can configure 802.3ad link aggregation 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

Module Requirements

For information about the modules that support 802.3ad link aggregation on ERX-14xx models, ERX-7xx models, and the ERX-310 router:

- See *ERX Module Guide, Table 1, Module Combinations* for detailed module specifications.
- See *ERX Module Guide, Appendix A, Module Protocol Support* for information about the modules that support 802.3ad link aggregation.

For information about the modules that support 802.3ad link aggregation on the E120 router and the E320 router:

- See *E120 and E320 Module Guide, Table 1, Modules and IOAs* for detailed module specifications.
- See *E120 and E320 Module Guide, Appendix A, IOA Protocol Support* for information about the modules that support 802.3ad link aggregation.

Interface Specifiers

The configuration task examples in this chapter use the format for ERX-7xx models, ERX-14xx models, and the ERX-310 router to specify 802.3ad link aggregation.

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

host1(config)#**interface gigabitEthernet 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*.

802.3ad Link Aggregation References

For more information about 802.3ad link aggregation implementations, consult the following resources:

- IEEE 802.1w (Rapid Reconfiguration of Spanning Tree)
- IEEE 802.3ad (Link Aggregation)

Configuring 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*.

Tasks to configure 802.3ad link aggregation interfaces are:

- Configuring an Ethernet Physical Interface on page 197
- Configuring a LAG Bundle on page 198
- Configuring IP for a LAG Bundle on page 198
- Configuring a VLAN Subinterface for a LAG Bundle on page 198
- Configuring a PPPoE Subinterface for a LAG Bundle on page 199
- Configuring MPLS for a LAG Bundle on page 199

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. Specify the MTU.

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

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

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

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 203.

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**

interface lag

- Use to create an IEEE 802.3ad LAG interface, also known as a 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.

mtu

- Use to specify the MTU for a LAG bundle.
- 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.

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.

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 201.

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 201.

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
```

Ethernet Link Redundancy Overview

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

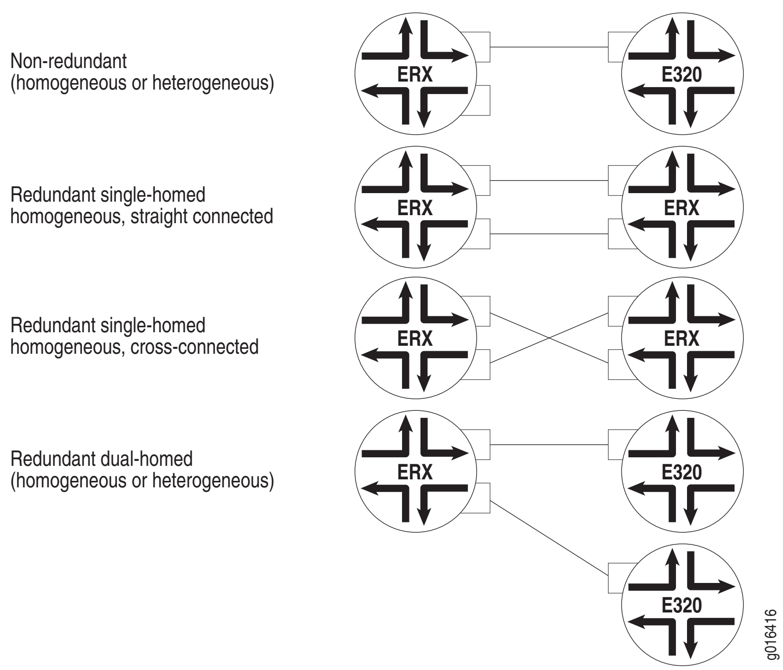

Ethernet Link Redundancy Configuration Diagrams

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

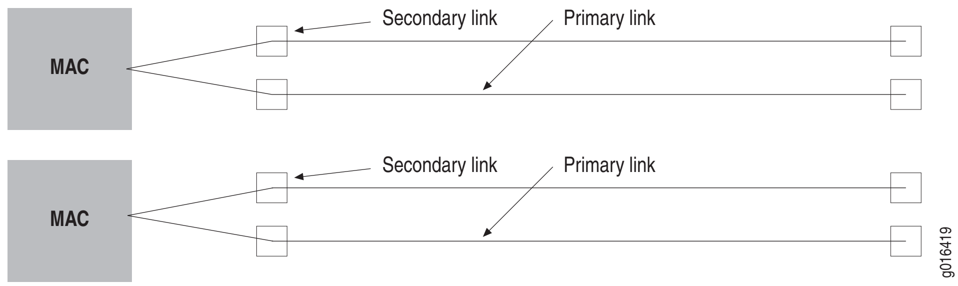

Figure 23 displays a single-homed configuration with port 0 backing up port 1 on a GE-2 line module.

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

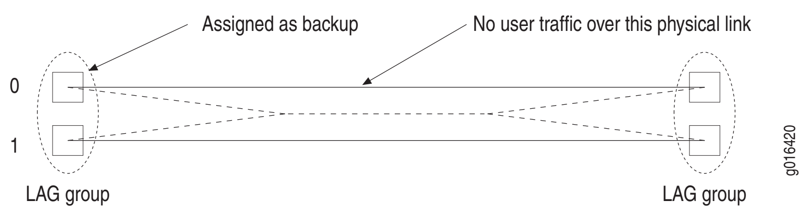

FE-8 Line Module Configurations

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)

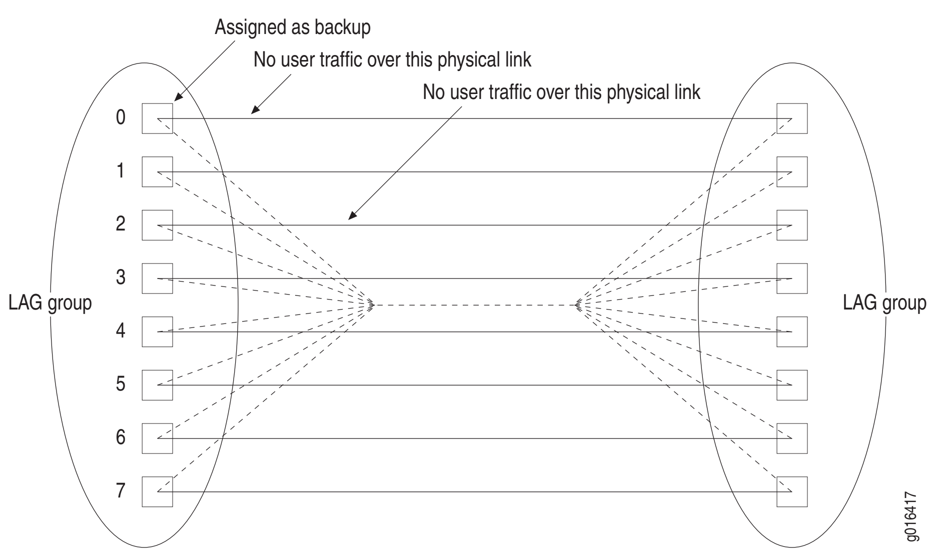

Figure 25 displays an FE-8 line module with four redundant Ethernet links in a 1:1 configuration.

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

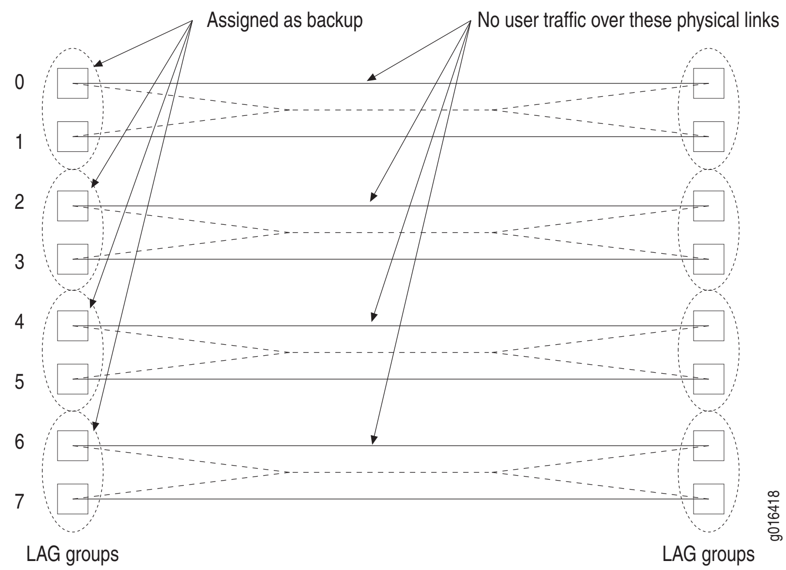

E120 and E320 Router Configurations

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)

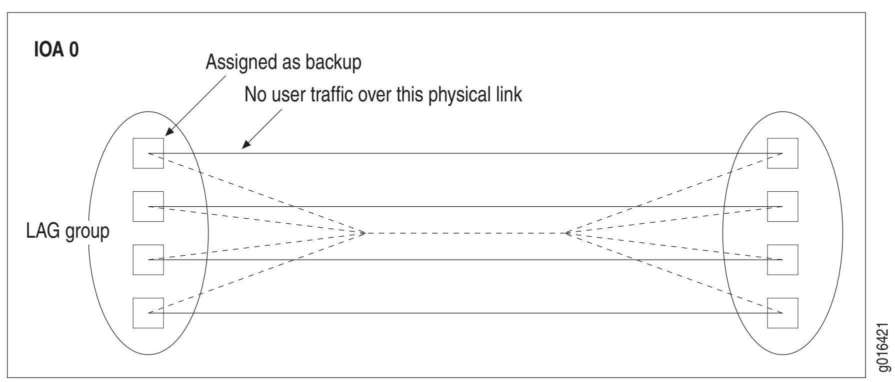

Figure 27 displays an E320 router with 1:N configuration across IOAs.

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

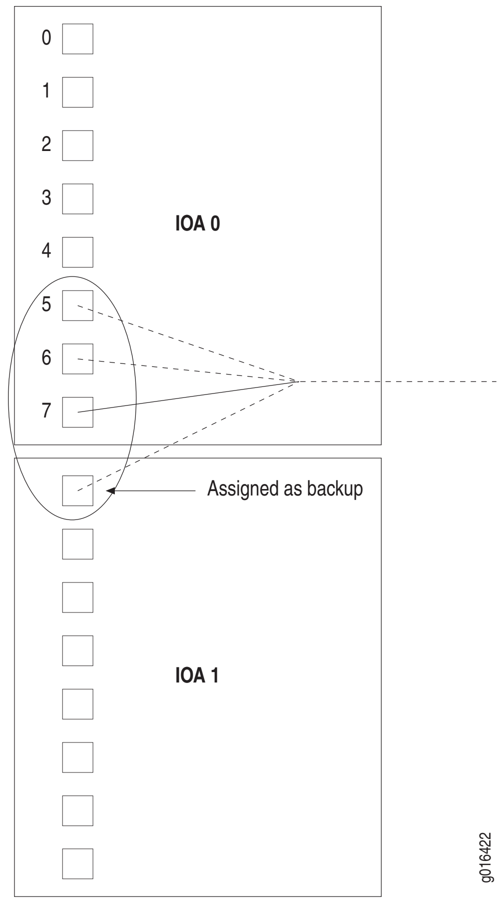

Dual-Homed Configurations with LAG Disabled

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)

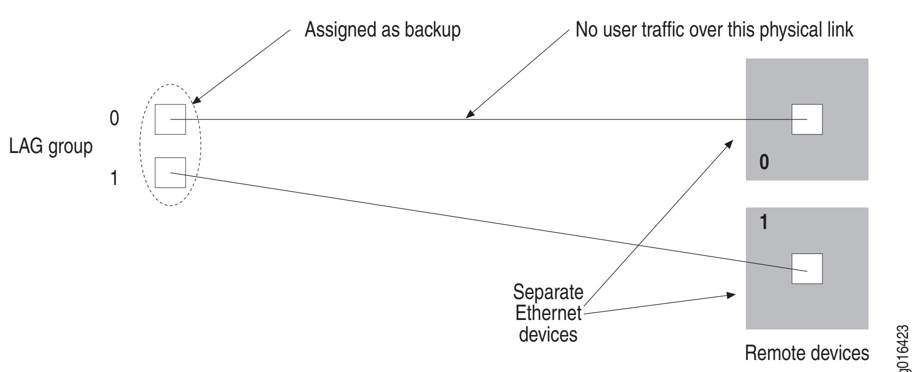

Ethernet Link Redundancy Behavior

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 194.

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. Disabling the transmitter also enables the remote end of the redundant link to be in the operational Down state, which might be a requirement for third-party equipment when supporting redundancy over LAG.

Enabling the transmitter provides for a quick LAG failover in the event one of the non-redundant links in the LAG fail. This is particularly true when LACP has been enabled on the LAG, because it can take several seconds for LACP to converge on a link. 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 212.

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

Table 11: 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

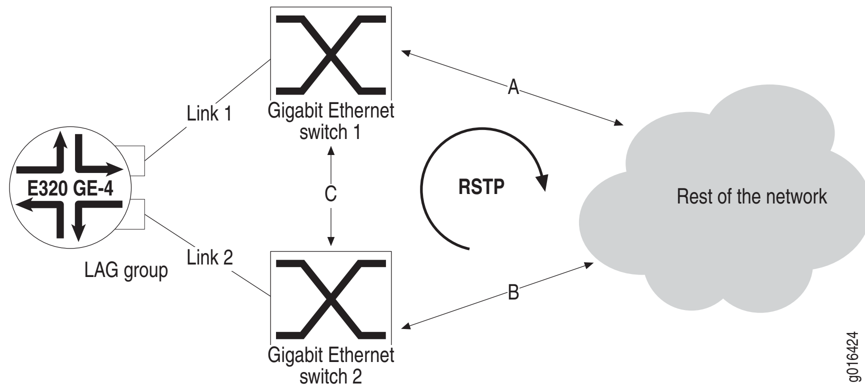

Figure 29 displays a network with RSTP enabled on Gigabit Ethernet switches 1 and 2. The local port receives bridge PDUs (BPDU), Ethernet broadcasts, and flooded unicast packets. If Link 1 is initially active and Link 2 is the backup, initial traffic destined for the LAG can be Ethernet broadcasts, PPPoE 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**
   host1(config-if)#**member-interface gigabitEthernet 1/1**

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.
  - Disabling the transmitter enables the remote end of the redundant link to also be in the operational Down state, which might be a requirement for third-party equipment when supporting redundancy over LAG.
  - Enabling the transmitter provides for a quick LAG failover in the event one of the non-redundant links in the LAG fail. This is particularly true when LACP has been enabled on the LAG, because it can take several seconds for LACP to converge on a link.
- 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.

Monitoring 802.3ad Link Aggregation

This section explains how to use the **show** commands to display the characteristics and the configured settings for 802.3ad link aggregation.

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*).

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
```