



Configuring EtherChannels

This chapter describes how to configure EtherChannel on the Layer 2 and Layer 3 interfaces of a Catalyst 3550 switch. EtherChannel provides fault-tolerant high-speed links between switches, routers, and servers. You can use it to increase the bandwidth between the wiring closets and the data center, and you can deploy it anywhere in the network where bottlenecks are likely to occur. EtherChannel provides automatic recovery for the loss of a link by redistributing the load across the remaining links. If a link fails, EtherChannel redirects traffic from the failed link to the remaining links in the channel without intervention.

This chapter consists of these sections:

- [Understanding EtherChannels, page 29-1](#)
- [Configuring EtherChannels, page 29-7](#)
- [Displaying EtherChannel, PAgP, and LACP Status, page 29-18](#)



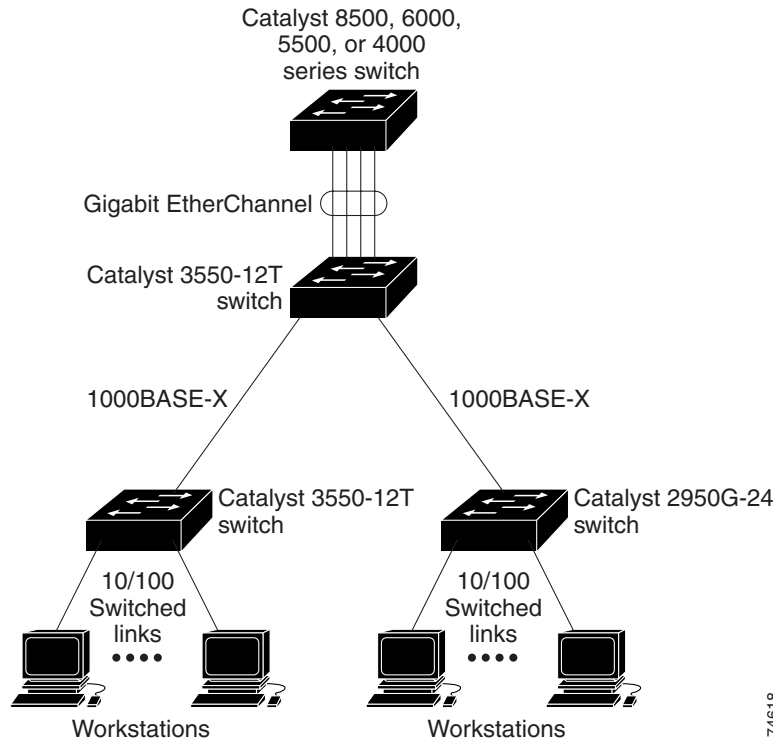
Note

For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release.

Understanding EtherChannels

An EtherChannel consists of individual Fast Ethernet or Gigabit Ethernet links bundled into a single logical link as shown in [Figure 29-1](#). The EtherChannel provides full-duplex bandwidth up to 800 Mbps (Fast EtherChannel) or 8 Gbps (Gigabit EtherChannel) between your switch and another switch or host.

Figure 29-1 Typical EtherChannel Configuration



Each EtherChannel can consist of up to eight compatibly configured Ethernet interfaces. All interfaces in each EtherChannel must be the same speed, and all must be configured as either Layer 2 or Layer 3 interfaces.

**Note**

The network device to which your switch is connected can impose its own limits on the number of interfaces in the EtherChannel. For Catalyst 3550 switches, the number of EtherChannels is limited to the number of ports of the same type.

If a link within an EtherChannel fails, traffic previously carried over that failed link changes to the remaining links within the EtherChannel. A trap is sent for a failure, identifying the switch, the EtherChannel, and the failed link. Inbound broadcast and multicast packets on one link in an EtherChannel are blocked from returning on any other link of the EtherChannel.


Understanding Port-Channel Interfaces

You create an EtherChannel for Layer 2 interfaces differently from Layer 3 interfaces. Both configurations involve logical interfaces.

- With Layer 3 interfaces, you manually create the logical interface by using the **interface port-channel** global configuration command.
- With Layer 2 interfaces, the logical interface is dynamically created.
- With both Layer 3 and 2 interfaces, you manually assign an interface to the EtherChannel by using the **channel-group** interface configuration command. This command binds the physical and logical ports together as shown in [Figure 29-2](#).

Each EtherChannel has a logical port-channel interface numbered from 1 to 64. The channel groups are also numbered from 1 to 64.

Figure 29-2 Relationship of Physical Ports, Logical Port Channels, and Channel Groups



When a port joins an EtherChannel, the physical interface for that port is shut down. When the port leaves the port-channel, its physical interface is brought up, and it has the same configuration as it had before joining the EtherChannel.



Note

Configuration changes made to the logical interface of an EtherChannel do not propagate to all the member ports of the channel.

Understanding the Port Aggregation Protocol and Link Aggregation Protocol

The Port Aggregation Protocol (PAgP) and Link Aggregation Control Protocol (LACP) facilitate the automatic creation of EtherChannels by exchanging packets between Ethernet interfaces. PAgP is a Cisco-proprietary protocol that can be run only on Cisco switches and on those switches licensed by licensed vendors to support PAgP. LACP is defined in IEEE 802.3AD and allows Cisco switches to manage Ethernet channels between switches that conform to the 802.3AD protocol.

By using one of these protocols, a switch learns the identity of partners capable of supporting either PAgP or LACP and learns the capabilities of each interface. It then dynamically groups similarly configured interfaces into a single logical link (channel or aggregate port); these interfaces are grouped based on hardware, administrative, and port parameter constraints. For example, PAgP groups the interfaces with the same speed, duplex mode, native VLAN, VLAN range, and trunking status and type. After grouping the links into an EtherChannel, PAgP adds the group to the spanning tree as a single switch port.

PAgP and LACP Modes

Table 29-1 shows the user-configurable EtherChannel modes for the **channel-group** interface configuration command. Switch interfaces exchange PAgP packets only with partner interfaces configured in the **auto** or **desirable** modes. Switch interfaces exchange LACP packets only with partner interfaces configured in the **active** or **passive** modes. Interfaces configured in the **on** mode do not exchange PAgP or LACP packets.

Table 29-1 EtherChannel Modes

| Mode | Description |
|------------------|---|
| active | Places an interface into an active negotiating state, in which the interface starts negotiations with other interfaces by sending LACP packets. |
| auto | Places an interface into a passive negotiating state, in which the interface responds to PAgP packets it receives but does not start PAgP packet negotiation. This setting minimizes the transmission of PAgP packets. |
| desirable | Places an interface into an active negotiating state, in which the interface starts negotiations with other interfaces by sending PAgP packets. |
| on | Forces the interface into an EtherChannel without PAgP or LACP. With the on mode, a usable EtherChannel exists only when an interface group in the on mode is connected to another interface group in the on mode. |
| passive | Places an interface into a passive negotiating state, in which the interface responds to LACP packets that it receives, but does not start LACP packet negotiation. This setting minimizes the transmission of LACP packets. |

Exchanging PAgP Packets

Both the **auto** and **desirable** PAgP modes allow interfaces to negotiate with partner interfaces to determine if they can form an EtherChannel based on criteria such as interface speed and, for Layer 2 EtherChannels, trunking state and VLAN numbers.

Interfaces can form an EtherChannel when they are in different PAgP modes as long as the modes are compatible. For example:

- An interface in the **desirable** mode can form an EtherChannel with another interface that is in the **desirable** or **auto** mode.
- An interface in the **auto** mode can form an EtherChannel with another interface in the **desirable** mode.

An interface in the **auto** mode cannot form an EtherChannel with another interface that is also in the **auto** mode because neither interface starts PAgP negotiation.

An interface in the **on** mode that is added to a port channel is forced to have the same characteristics as the already existing **on** mode interfaces in the channel.

If your switch is connected to a partner that is PAgP-capable, you can configure the switch interface for nonsilent operation by using the **non-silent** keyword. If you do not specify **non-silent** with the **auto** or **desirable** mode, silent mode is assumed.

The silent mode is used when the switch is connected to a device that is not PAgP-capable and seldom, if ever, sends packets. An example of a silent partner is a file server or a packet analyzer that is not generating traffic. In this case, running PAgP on a physical port connected to a silent partner prevents that switch port from ever becoming operational; however, the silent setting allows PAgP to operate, to attach the interface to a channel group, and to use the interface for transmission.

**Note**

An Etherchannel cannot be configured in both the PAgP and LACP modes.

Exchanging LACP Packets

Both the **active** and **passive** LACP modes allow interfaces to negotiate with partner interfaces to determine if they can form an EtherChannel based on criteria such as interface speed and, for Layer 2 EtherChannels, trunking state and VLAN numbers.

Interfaces can form an EtherChannel when they are in different LACP modes as long as the modes are compatible. For example:

- An interface in the **active** mode can form an EtherChannel with another interface that is in the **active** or **passive** mode.
- An interface in the **active** mode can form an EtherChannel with another interface in the **passive** mode.

An interface in the **passive** mode cannot form an EtherChannel with another interface that is also in the **passive** mode because neither interface starts LACP negotiation.

An interface in the **on** mode that is added to a port channel is forced to have the same characteristics as the already existing **on** mode interfaces in the channel.

**Note**

An Etherchannel cannot be configured in both the PAgP and LACP modes.

**Caution**

You should exercise care when setting the mode to **on** (manual configuration). All ports configured in the **on** mode are bundled in the same group and are forced to have similar characteristics. If the group is misconfigured, packet loss or spanning-tree loops might occur.

Physical Learners and Aggregate-Port Learners

Network devices are classified as PAgP physical learners or aggregate-port learners. A device is a physical learner if it learns addresses by physical ports and directs transmissions based on that knowledge. A device is an aggregate-port learner if it learns addresses by aggregate (logical) ports.

When a device and its partner are both aggregate-port learners, they learn the address on the logical port-channel. The device sends packets to the source by using any of the interfaces in the EtherChannel. With aggregate-port learning, it is not important on which physical port the packet arrives.

PAgP cannot automatically detect when the partner device is a physical learner and the local device is an aggregate-port learner. Therefore, you must manually set the learning method on the local device or source-based distribution by using the **pagp learn-method** interface configuration command. With source-based distribution, any given source MAC address is always sent on the same physical port.

You can also configure a single interface within the group for all transmissions and use other interfaces for hot standby. The unused interfaces in the group can be swapped into operation in just a few seconds if the selected single interface loses hardware-signal detection. You can configure which interface is always selected for packet transmission by changing its priority by using the **pagp port-priority** interface configuration command. The higher the priority, the more likely that the port will be selected.

PAgP and LACP Interaction with Other Features

The Dynamic Trunking Protocol (DTP) and Cisco Discovery Protocol (CDP) send and receive packets over the physical interfaces in the EtherChannel. Trunk ports send and receive PAgP and LACP protocol data units (PDUs) on the lowest numbered VLAN.

Spanning tree sends packets over the first interface in the EtherChannel.

The MAC address of a Layer 3 EtherChannel is the MAC address of the first interface in the port-channel.

PAgP sends and receives PAgP PDUs only from interfaces that have PAgP enabled for the auto or desirable mode. LACP sends and receives LACP PDUs only from interfaces that have LACP enabled for the active or passive mode.

Understanding Load Balancing and Forwarding Methods

EtherChannel balances the traffic load across the links in a channel by randomly associating a newly-learned MAC address with one of the links in the channel.

With source-MAC address forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the channel based on the source-MAC address of the incoming packet. Therefore, to provide load balancing, packets from different hosts use different ports in the channel, but packets from the same host use the same port in the channel (and the MAC address learned by the switch does not change).

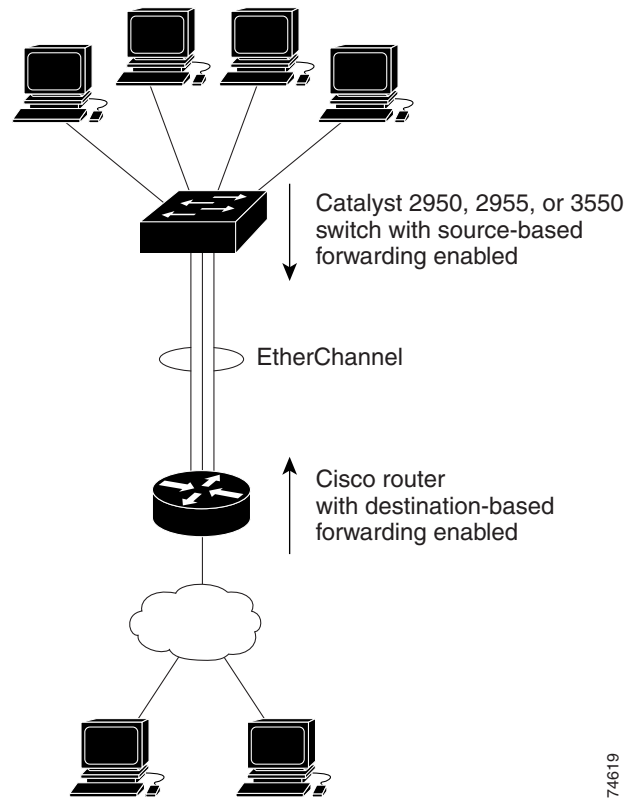
When source-MAC address forwarding is used, load distribution based on the source and destination IP address is also enabled for routed IP traffic. All routed IP traffic chooses a port based on the source and destination IP address. Packets between two IP hosts always use the same port in the channel, and traffic between any other pair of hosts can use a different port in the channel.

With destination-MAC address forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the channel based on the destination host's MAC address of the incoming packet. Therefore, packets to the same destination are forwarded over the same port, and packets to a different destination are sent on a different port in the channel. You configure the load balancing and forwarding method by using the **port-channel load-balance** global configuration command.

In [Figure 29-3](#), multiple workstations are connected to a switch, and an EtherChannel connects the switch to the router. Source-based load balancing is used on the switch end of the EtherChannel to ensure that the switch efficiently uses the bandwidth of the router by distributing traffic from the workstation across the physical links. Since the router is a single MAC address device, it uses destination-based load balancing to efficiently spread the traffic to the workstations across the physical links in the EtherChannel.

Use the option that provides the greatest variety in your configuration. For example, if the traffic on a channel is going only to a single MAC address, using the destination-MAC address always chooses the same link in the channel; using source addresses or IP addresses might result in better load balancing.

Figure 29-3 Load Distribution and Forwarding Methods



74619

Configuring EtherChannels

These sections describe how to configure EtherChannel on Layer 2 and Layer 3 interfaces:

- [Default EtherChannel Configuration, page 29-8](#)
- [EtherChannel Configuration Guidelines, page 29-8](#)
- [Configuring Layer 2 EtherChannels, page 29-9](#)
- [Configuring Layer 3 EtherChannels, page 29-11](#)
- [Configuring EtherChannel Load Balancing, page 29-14](#)
- [Configuring the PAgP Learn Method and Priority, page 29-15](#)



Note

Make sure that the interfaces are correctly configured (see the “[EtherChannel Configuration Guidelines](#)” section on page 29-8).



Note

After you configure an EtherChannel, configuration changes applied to the port-channel interface apply to all the physical interfaces assigned to the port-channel interface, and configuration changes applied to the physical interface affect only the interface where you apply the configuration.

Default EtherChannel Configuration

Table 29-2 shows the default EtherChannel configuration.

Table 29-2 Default EtherChannel Configuration

| Feature | Default Setting |
|--|--|
| Channel groups | None assigned. |
| Layer 3 port-channel logical interface | None defined. |
| PAgP mode | No default. |
| PAgP learn method | Aggregate-port learning on all interfaces. |
| PAgP priority | 128 on all interfaces. |
| LACP learn method | Aggregate-port learning on all interfaces. |
| LACP priority | 32768 on all interfaces. |
| Load balancing | Load distribution on the switch is based on the source-MAC address of the incoming packet. Load distribution based on the source and destination IP address is also enabled for routed IP traffic. |

EtherChannel Configuration Guidelines

If improperly configured, some EtherChannel interfaces are automatically disabled to avoid network loops and other problems. Follow these guidelines to avoid configuration problems:

- Configure an EtherChannel with up to eight Ethernet interfaces of the same type.



Note Do not configure a GigaStack GBIC port as part of an EtherChannel.

- Configure all interfaces in an EtherChannel to operate at the same speeds and duplex modes.
- Enable all interfaces in an EtherChannel. An interface in an EtherChannel that is disabled by using the **shutdown** interface configuration command is treated as a link failure, and its traffic is transferred to one of the remaining interfaces in the EtherChannel.
- When a group is first created, all ports follow the parameters set for the first port to be added to the group. If you change the configuration of one of these parameters, you must also make the changes to all ports in the group:
 - Allowed-VLAN list
 - Spanning-tree path cost for each VLAN
 - Spanning-tree port priority for each VLAN
 - Spanning-tree Port Fast setting
- An EtherChannel interface that is configured as a Switched Port Analyzer (SPAN) destination port does not join the group until it is deconfigured as a SPAN destination port.
- Do not configure a secure port as part of an EtherChannel.

- Before enabling 802.1X on the port, you must first remove it from the EtherChannel. If you try to enable 802.1X on an EtherChannel or on an active port in an EtherChannel, an error message appears, and 802.1X is not enabled. If you enable 802.1X on a not-yet active port of an EtherChannel, the port does not join the EtherChannel.
- For Layer 2 EtherChannels:
 - Assign all interfaces in the EtherChannel to the same VLAN, or configure them as trunks. Interfaces with different native VLANs cannot form an EtherChannel.
 - If you configure an EtherChannel from trunk interfaces, verify that the trunking mode (ISL or 802.1Q) is the same on all the trunks. Inconsistent trunk modes on EtherChannel interfaces can have unexpected results.
 - An EtherChannel supports the same allowed range of VLANs on all the interfaces in a trunking Layer 2 EtherChannel. If the allowed range of VLANs is not the same, the interfaces do not form an EtherChannel even when PAgP is set to the **auto** or **desirable** mode.
 - Interfaces with different spanning-tree path costs can form an EtherChannel if they are otherwise compatibly configured. Setting different spanning-tree path costs does not, by itself, make interfaces incompatible for the formation of an EtherChannel.
- For Layer 3 EtherChannels, assign the Layer 3 address to the port-channel logical interface, not to the physical interfaces in the channel.

Configuring Layer 2 EtherChannels

You configure Layer 2 EtherChannels by configuring the Ethernet interfaces with the **channel-group** interface configuration command, which creates the port-channel logical interface. You cannot put a Layer 2 interface into a manually created port-channel interface.



Note

Layer 2 interfaces must be connected and functioning for Cisco the software to create port-channel interfaces.

Beginning in privileged EXEC mode, follow these steps to assign a Layer 2 Ethernet interface to a Layer 2 EtherChannel:

| | Command | Purpose |
|--------|---|--|
| Step 1 | configure terminal | Enter global configuration mode. |
| Step 2 | interface <i>interface-id</i> | Enter interface configuration mode, and specify a physical interface to configure. Valid interfaces include physical interfaces. Up to eight interfaces of the same type and speed can be configured for the same group. |
| Step 3 | switchport mode {access trunk} switchport access vlan <i>vlan-id</i> | Assign all interfaces as static-access ports in the same VLAN, or configure them as trunks. If you configure the interface as a static-access port, assign it to only one VLAN. The range is 1 to 4094. |

| | Command | Purpose |
|--------|--|--|
| Step 4 | channel-group <i>channel-group-number</i> mode { { auto [non-silent] desirable [non-silent] on } { active passive } } | <p>Assign the interface to a channel group, and specify the PAgP or LACP mode.</p> <p>For <i>channel-group-number</i>, the range is 1 to 64. Each EtherChannel can have up to eight compatibly configured Ethernet interfaces.</p> <p>For mode, select one of these keywords:</p> <ul style="list-style-type: none"> • active—Enables LACP only if an LACP device is detected. It places an interface into an active negotiating state, in which the interface starts negotiations with other interfaces by sending LACP packets. • auto—Enables PAgP only if a PAgP device is detected. It places an interface into a passive negotiating state, in which the interface responds to PAgP packets it receives but does not start PAgP packet negotiation. • desirable—Unconditionally enables PAgP. It places an interface into an active negotiating state, in which the interface starts negotiations with other interfaces by sending PAgP packets. • on—Forces the interface to channel without PAgP. With the on mode, a usable EtherChannel exists only when an interface group in the on mode is connected to another interface group in the on mode. • non-silent—If your switch is connected to a partner that is PAgP-capable, you can configure the switch interface for nonsilent operation. You can configure an interface with the non-silent keyword for use with the auto or desirable mode. If you do not specify non-silent with the auto or desirable mode, silent is assumed. The silent setting is for connections to file servers or packet analyzers. This setting allows PAgP to operate, to attach the interface to a channel group, and to use the interface for transmission. • passive—Enables LACP on an interface and places it into a passive negotiating state, in which the interface responds to LACP packets that it receives, but does not start LACP packet negotiation. <p>For information on compatible PAgP and LACP modes for the switch and its partner, see the “PAgP and LACP Modes” section on page 29-4.</p> |
| Step 5 | end | Return to privileged EXEC mode. |
| Step 6 | show running-config | Verify your entries. |
| Step 7 | copy running-config startup-config | (Optional) Save your entries in the configuration file. |

To remove an interface from the EtherChannel group, use the **no channel-group** interface configuration command.

This example shows how to assign Gigabit Ethernet interfaces 0/4 and 0/5 as static-access ports in VLAN 10 to channel 5 with the PAgP mode **desirable**:

```
Switch# configure terminal
Switch(config)# interface range gigabitethernet0/4 -5
Switch(config-if-range)# switchport mode access
Switch(config-if-range)# switchport access vlan 10
Switch(config-if-range)# channel-group 5 mode desirable
Switch(config-if-range)# end
```

Configuring Layer 3 EtherChannels

To configure Layer 3 EtherChannels, you create the port-channel logical interface and then put the Ethernet interfaces into the port-channel as described in the next two sections.

Creating Port-Channel Logical Interfaces

When configuring Layer 3 EtherChannels, you must manually create the port-channel logical interface first by using the **interface port-channel** global configuration command. Then, you put the logical interface into the channel group by using the **channel-group** interface configuration command.



Note

To move an IP address from a physical interface to an EtherChannel, you must delete the IP address from the physical interface before configuring it on the port-channel interface.

Beginning in privileged EXEC mode, follow these steps to create a port-channel interface for a Layer 3 EtherChannel:

| | Command | Purpose |
|--------|---|--|
| Step 1 | configure terminal | Enter global configuration mode. |
| Step 2 | interface port-channel <i>port-channel-number</i> | Enter interface configuration mode, and create the port-channel logical interface. For <i>port-channel-number</i> , the range is 1 to 64. |
| Step 3 | no switchport | Put the interface into Layer 3 mode. |
| Step 4 | ip address <i>ip-address mask</i> | Assign an IP address and subnet mask to the EtherChannel. |
| Step 5 | end | Return to privileged EXEC mode. |
| Step 6 | show etherchannel <i>channel-group-number detail</i> | Verify your entries. |
| Step 7 | copy running-config startup-config | (Optional) Save your entries in the configuration file. |
| Step 8 | | Assign an Ethernet interface to the Layer 3 EtherChannel. For more information, see the “Configuring the Physical Interfaces” section on page 29-12. |

To remove the port-channel, use the **no interface port-channel** *port-channel-number* global configuration command.

This example shows how to create the logical port channel (5) and assign 172.10.20.10 as its IP address:

```
Switch# configure terminal
Switch(config)# interface port-channel 5
Switch(config-if)# no switchport
```

```
Switch(config-if)# ip address 172.10.20.10 255.255.255.0
Switch(config-if)# end
```

Configuring the Physical Interfaces

Beginning in privileged EXEC mode, follow these steps to assign an Ethernet interface to a Layer 3 EtherChannel:

| | Command | Purpose |
|--------|--------------------------------------|--|
| Step 1 | configure terminal | Enter global configuration mode. |
| Step 2 | interface <i>interface-id</i> | Enter interface configuration mode, and specify a physical interface to configure. Valid interfaces include physical interfaces. Up to eight interfaces of the same type and speed can be configured for the same group. |
| Step 3 | no ip address | Ensure that there is no IP address assigned to the physical interface. |

| | Command | Purpose |
|--------|---|--|
| Step 4 | channel-group <i>channel-group-number</i> mode { auto [non-silent] desirable [non-silent] on active passive } | Assign the interface to a channel group, and specify the PAgP or LACP mode. For <i>channel-group-number</i> , the range is 1 to 64. This number must be the same as the <i>port-channel-number</i> (logical port) configured in the “Creating Port-Channel Logical Interfaces” section on page 29-11. Each EtherChannel can consist of up to eight compatibly configured Ethernet interfaces. For mode , select one of these keywords: <ul style="list-style-type: none"> • active—Enables LACP only if an LACP device is detected. It places an interface into an active negotiating state, in which the interface starts negotiations with other interfaces by sending LACP packets. • auto—Enables PAgP only if a PAgP device is detected. It places an interface into a passive negotiating state, in which the interface responds to PAgP packets it receives but does not start PAgP packet negotiation. • desirable—Unconditionally enables PAgP. It places an interface into an active negotiating state, in which the interface starts negotiations with other interfaces by sending PAgP packets. • non-silent—If your switch is connected to a partner that is PAgP capable, you can configure the switch interface for nonsilent operation. You can configure an interface with the non-silent keyword for use with the auto or desirable mode. If you do not specify non-silent with the auto or desirable mode, silent is assumed. The silent setting is for connections to file servers or packet analyzers; this setting allows PAgP to operate, to attach the interface to a channel group, and to use the interface for transmission. • on—Forces the interface to channel without PAgP or LACP. With the on mode, a usable EtherChannel exists only when an interface group in the on mode is connected to another interface group in the on mode. • passive—Enables LACP on an interface and places it into a passive negotiating state, in which the interface responds to LACP packets that it receives, but does not start LACP packet negotiation. For information on compatible PAgP modes for the switch and its partner, see the “PAgP and LACP Modes” section on page 29-4. |
| Step 5 | end | Return to privileged EXEC mode. |
| Step 6 | show running-config | Verify your entries. |
| Step 7 | copy running-config startup-config | (Optional) Save your entries in the configuration file. |

To remove an interface from the EtherChannel group, use the **no channel-group** interface configuration command.

This example shows how to assign Gigabit Ethernet interfaces 0/4 and 0/5 to channel 5 with the PAgP mode **desirable**:

```
Switch# configure terminal
Switch(config)# interface range gigabitethernet0/4 -5
Switch(config-if-range)# no ip address
Switch(config-if-range)# channel-group 5 mode desirable
Switch(config-if-range)# end
```

Configuring EtherChannel Load Balancing

This section describes how to configure EtherChannel load balancing by using source-based or destination-based forwarding methods. For more information, see the [“Understanding Load Balancing and Forwarding Methods”](#) section on page 29-6.

Beginning in privileged EXEC mode, follow these steps to configure EtherChannel load balancing:

| | Command | Purpose |
|--------|--|--|
| Step 1 | configure terminal | Enter global configuration mode. |
| Step 2 | port-channel load-balance {dst-mac src-mac} | <p>Configure an EtherChannel load-balancing method.</p> <p>The default is src-mac.</p> <p>Select one of these keywords to determine the load-distribution method:</p> <ul style="list-style-type: none"> • dst-mac—Load distribution is based on the destination-host MAC address of the incoming packet. Packets to the same destination are sent on the same port, but packets to different destinations are sent on different ports in the channel. • src-mac—Load distribution is based on the source-MAC address of the incoming packet. Packets from different hosts use different ports in the channel, but packets from the same host use the same port in the channel. <p>When src-mac is used, load distribution based on the source and destination IP address is also enabled. For all IP traffic being routed, the switch chooses a port for transmission based on the source and destination IP address. Packets between two IP hosts always use the same port for packet transmission, but packets between any other pair of hosts might use a different transmission port.</p> |
| Step 3 | end | Return to privileged EXEC mode. |
| Step 4 | show etherchannel load-balance | Verify your entries. |
| Step 5 | copy running-config startup-config | (Optional) Save your entries in the configuration file. |

To return EtherChannel load balancing to the default configuration, use the **no port-channel load-balance** global configuration command.

Configuring the PAgP Learn Method and Priority

Network devices are classified as PAgP physical learners or aggregate-port learners. A device is a physical learner if it learns addresses by physical ports and directs transmissions based on that knowledge. A device is an aggregate-port learner if it learns addresses by aggregate ports.

For compatibility with Catalyst 1900 series switches, configure the PAgP learning method on the Catalyst 3550 switches to learn source-MAC addresses on the physical port. The switch then sends packets to the Catalyst 1900 switch using the same interface in the EtherChannel from which it learned the source address.



Note

The Catalyst 3550 supports address learning only on aggregate ports even though the **physical-port** keyword is provided in the CLI. The **pagp learn-method** command and the **pagp port-priority** command have no effect on the switch hardware, but they are required for PAgP interoperability with devices that only support address learning by physical ports, such as the Catalyst 1900 switch.

When the link partner to the Catalyst 3550 switch is a physical learner, we recommend that you configure the switch as a physical-port learner by using the **pagp learn-method physical-port** interface configuration command and to set the load-distribution method based on the source MAC address by using the **port-channel load-balance src-mac** global configuration command. Use the **pagp learn-method** command only in this situation.

Beginning in privileged EXEC mode, follow these steps to configure your switch as a PAgP physical-port learner and to adjust the priority so that the same port in the bundle is selected for sending packets:

| | Command | Purpose |
|--------|---|--|
| Step 1 | configure terminal | Enter global configuration mode. |
| Step 2 | interface <i>interface-id</i> | Enter interface configuration mode, and specify the interface for transmission. |
| Step 3 | pagp learn-method physical-port | Select the PAgP learning method. By default, aggregation-port learning is selected, which means the switch sends packets to the source by using any of the interfaces in the EtherChannel. With aggregate-port learning, it is not important on which physical port the packet arrives. Select physical-port to connect with another switch that is a physical learner. Make sure to configure the port-channel load-balance global configuration command to src-mac as described in the “Configuring EtherChannel Load Balancing” section on page 29-14. The learning method must be configured the same at both ends of the link. |
| Step 4 | pagp port-priority <i>priority</i> | Assign a priority so that the selected interface is chosen for packet transmission. For <i>priority</i> , the range is 0 to 255. The default is 128. The higher the priority, the more likely that the interface will be used for PAgP transmission. |
| Step 5 | end | Return to privileged EXEC mode. |

| | Command | Purpose |
|--------|--|---|
| Step 6 | show running-config or show pagp channel-group-number internal | Verify your entries. |
| Step 7 | copy running-config startup-config | (Optional) Save your entries in the configuration file. |

To return the priority to its default setting, use the **no pagp port-priority** interface configuration command. To return the learning method to its default setting, use the **no pagp learn-method** interface configuration command.

Configuring the LACP Port Priority

You can set the priority for each port in an EtherChannel that is configured for LACP by using the **lacp port-priority** privileged EXEC command. The range is from 1 to 65535. Beginning in privileged EXEC mode, follow these steps to configure the LACP port priority:

| | Command | Purpose |
|--------|--|---|
| Step 1 | configure terminal | Enter global configuration mode. |
| Step 2 | interface interface-id | Enter interface configuration mode, and specify the interface for transmission. |
| Step 3 | lacp port-priority priority-value | Select the LACP port priority value. For priority-value, the range is 1 to 65535. By default, the priority value is 32768. The lower the range, the more likely that the interface will be used for LACP transmission. |
| Step 4 | end | Return to privileged EXEC mode. |
| Step 5 | show running-config or show lacp channel-group-number internal | Verify your entries. |
| Step 6 | copy running-config startup-config | (Optional) Save your entries in the configuration file. |

Configuring Hot Standby Ports

When enabled, LACP tries to configure the maximum number of LACP-compatible ports in a channel, up to a maximum of 16 ports. Only eight LACP links can be active at one time. Any additional links are put in a hot standby state. If one of the active links becomes inactive, a link that is in hot standby mode becomes active in its place.

If more than eight links are configured for an EtherChannel group, the software determines which of the hot standby ports to make active based on:

- LACP port-priority
- Port ID

All ports default to the same port priority. You can change the port priority of LACP EtherChannel ports to specify which hot standby links become active first by using the **lacp port-priority** interface configuration command to set the port priority to a value lower than the default of 32768.

The hot standby ports that have lower port numbers become active in the channel first unless the port priority is configured to be a lower number than the default value of 32768.

**Note**

If LACP is not able to aggregate all the ports that are compatible (for example, the remote system might have more restrictive hardware limitations), all the ports that cannot be actively included in the EtherChannel are put in hot standby state and are used only if one of the channeled ports fails.

Configuring the LACP System Priority

You can set the system priority for all of the EtherChannels that are configured for LACP by using the **lacp system-priority** privileged EXEC command. The range is from 1 to 65535.

**Note**

The **lacp system-priority** command is global. You cannot set a system priority for each LACP-configured channel separately.

We recommend using this command only when there are a combination of LACP-configured EtherChannels that are in both **active** and **standby** modes.

Beginning in privileged EXEC mode, follow these steps to configure the LACP system priority:

| | Command | Purpose |
|--------|--|---|
| Step 1 | configure terminal | Enter global configuration mode. |
| Step 2 | lacp system-priority <i>priority-value</i> | Select the LACP system priority value. For <i>priority-value</i> , the range is 1 to 65535. By default, the priority value is 32768. The lower the range, the higher the system priority. The switch with the lower system priority value determines which links between LACP partner switches are active and which are in standby for each LACP EtherChannel. |
| Step 3 | end | Return to privileged EXEC mode. |
| Step 4 | show running-config or show lacp <i>channel-group-number</i> internal | Verify your entries. |
| Step 5 | copy running-config startup-config | (Optional) Save your entries in the configuration file. |

Displaying EtherChannel, PAgP, and LACP Status

You can use the privileged EXEC commands described in [Table 29-3](#) to display EtherChannel, PAgP, and LACP status information:

Table 29-3 *Commands for Displaying EtherChannel, PAgP, and LACP Status*

| Command | Description |
|---|---|
| show etherchannel [<i>channel-group-number</i>] { detail load-balance port port-channel summary } | Displays EtherChannel information in a detailed and one-line summary form. Also displays the load-balance or frame-distribution scheme, port, and port-channel information. |
| show pagp [<i>channel-group-number</i>] { counters internal neighbor } ¹ | Displays PAgP information such as traffic information, the internal PAgP configuration, and neighbor information. |
| show lacp [<i>channel-group-number</i>] { counters internal neighbor } ² | Displays LACP information such as traffic information, the internal PAgP configuration, and neighbor information. |

1. You can clear PAgP channel-group information and traffic filters by using the **clear pagp** [*channel-group-number*] [**counters**] | **counters**) privileged EXEC command.
2. You can clear LACP channel-group information and traffic filters by using the **clear lacp** [*channel-group-number*] [**counters**] | **counters**) privileged EXEC command.

For detailed information about the fields in the command outputs, refer to the command reference for this release.