



Configuring EtherChannels

This chapter describes how to configure EtherChannels on Layer 2 and Layer 3 interfaces on the Catalyst 3750 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. Unless otherwise noted, the term *switch* refers to a standalone switch and to a switch stack.



Note

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

This chapter consists of these sections:

- [Understanding EtherChannels, page 25-1](#)
- [Configuring EtherChannels, page 25-9](#)
- [Displaying EtherChannel and PAgP Status, page 25-19](#)

Understanding EtherChannels


These sections describe how EtherChannels work:

- [EtherChannel Overview, page 25-2](#)
- [Port-Channel Interfaces, page 25-4](#)
- [Port Aggregation Protocol, page 25-5](#)
- [Load Balancing and Forwarding Methods, page 25-6](#)
- [EtherChannel and Switch Stacks, page 25-8](#)

EtherChannel Overview

An EtherChannel consists of individual Fast Ethernet or Gigabit Ethernet links bundled into a single logical link as shown in [Figure 25-1](#).

Figure 25-1 Typical EtherChannel Configuration



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.

Each EtherChannel can consist of up to eight compatibly configured Ethernet interfaces. All interfaces in each EtherChannel must be configured as either Layer 2 or Layer 3 interfaces. For Catalyst 3750 switches, the number of EtherChannels is limited to 12. For more information, see the [“EtherChannel Configuration Guidelines”](#) section on page 25-10. The EtherChannel Layer 3 interfaces are made up of routed ports. Routed ports are physical ports configured to be in Layer 3 mode by using the **no switchport** interface configuration command. For more information, see the [Chapter 9, “Configuring Interface Characteristics.”](#)

You can create an EtherChannel on a standalone switch, on a single switch in the stack, or on multiple switches in the stack (known as cross-stack EtherChannel). See [Figure 25-2](#) and [Figure 25-3](#).

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.

Figure 25-2 Single-Switch EtherChannel





Figure 25-3 Cross-Stack EtherChannel



Port-Channel Interfaces

When you create an EtherChannel, a port-channel logical interface is involved:

- With Layer 2 interfaces, use the **channel-group** interface configuration command to dynamically create the port-channel logical interface.


You also can use the **interface port-channel** *port-channel-number* global configuration command to manually create the port-channel logical interface, but then you must use the **channel-group** *channel-group-number* command to bind the logical interface to a physical port. The *channel-group-number* can be the same as the *port-channel-number*, or you can use a new number. If you use a new number, the **channel-group** command dynamically creates a new port channel.

- With Layer 3 interfaces, you should manually create the logical interface by using the **interface port-channel** global configuration command followed by the **no switchport** interface configuration command. Then you manually assign an interface to the EtherChannel by using the **channel-group** interface configuration command.

For both Layer 2 and Layer 3 interfaces, the **channel-group** command binds the physical and logical ports together as shown in Figure 25-4.

Each EtherChannel has a port-channel logical interface numbered from 1 to 12. This port-channel interface number corresponds to the one specified with the **channel-group** interface configuration command.

Figure 25-4 Relationship of Physical Ports, Logical Port Channels, and Channel Groups



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. Configuration changes applied to the physical interface affect only the interface where you apply the configuration. To change the parameters of all ports in an EtherChannel, apply configuration commands to the port-channel interface, for example, spanning-tree commands or commands to configure a Layer 2 EtherChannel as a trunk.

Port Aggregation Protocol

The Port Aggregation Protocol (PAgP) is a Cisco-proprietary protocol that can be run only on Cisco switches and on those switches licensed by vendors to support PAgP. PAgP facilitates the automatic creation of EtherChannels by exchanging PAgP packets between Ethernet interfaces. You can use PAgP only in single-switch EtherChannel configurations; PAgP cannot be enabled on cross-stack EtherChannels. For more information, see the [“EtherChannel Configuration Guidelines” section on page 25-10](#).

By using PAgP, the switch stack learns the identity of partners capable of supporting PAgP and the capabilities of each interface. It then dynamically groups similarly configured interfaces (on a single switch in the stack) into a single logical link (channel or aggregate port). Similarly configured 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 Modes

Table 25-1 shows the user-configurable EtherChannel PAgP modes for the **channel-group** interface configuration command.

Table 25-1 EtherChannel PAgP Modes

Mode	Description
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 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. This is the only setting that is supported when the EtherChannel members are from different switches in the switch stack (cross-stack EtherChannel).

Switch interfaces exchange PAgP packets only with partner interfaces configured in the **auto** or **desirable** modes. Interfaces configured in the **on** mode do not exchange PAgP packets.

Both the **auto** and **desirable** 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.

**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.

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.

Use the silent mode 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.

PAgP Interaction with Other Features

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

In Layer 2 EtherChannels, the first port in the channel that comes up provides its MAC address to the EtherChannel. If this port is removed from the bundle, one of the remaining ports in the bundle provides its MAC address to the EtherChannel. For Layer 3 EtherChannels, the MAC address is allocated by the stack master as soon as the interface is created (through the **interface port-channel** global configuration command).

PAgP sends and receives PAgP PDUs only from interfaces that are up and have PAgP enabled for the auto or desirable mode.

Load Balancing and Forwarding Methods

EtherChannel balances the traffic load across the links in a channel by reducing part of the binary pattern formed from the addresses in the frame to a numerical value that selects one of the links in the channel. EtherChannel load balancing can use MAC addresses or IP addresses, source or destination addresses, or both source and destination addresses. The selected mode applies to all EtherChannels configured on the switch. You configure the load balancing and forwarding method by using the **port-channel load-balance** global configuration command.

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.

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.

With source-and-destination MAC address forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the channel based on both the source and destination MAC addresses. This forwarding method, a combination source-MAC and destination-MAC address

forwarding methods of load distribution, can be used if it is not clear whether source-MAC or destination-MAC address forwarding is better suited on a particular switch. With source-and-destination MAC address forwarding, packets sent from host A to host B, host A to host C, and host C to host B could all use different ports in the channel.

With source-IP address-based forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the EtherChannel based on the source-IP address of the incoming packet. Therefore, to provide load-balancing, packets from different IP addresses use different ports in the channel, but packets from the same IP address use the same port in the channel.


With destination-IP address-based forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the EtherChannel based on the destination-IP address of the incoming packet. Therefore, to provide load-balancing, packets from the same IP source address sent to different IP destination addresses could be sent on different ports in the channel. But packets sent from different source IP addresses to the same destination IP address are always sent on the same port in the channel.

With source-and-destination IP address-based forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the EtherChannel based on both the source and destination IP addresses of the incoming packet. This forwarding method, a combination of source-IP and destination-IP address-based forwarding, can be used if it is not clear whether source-IP or destination-IP address-based forwarding is better suited on a particular switch. In this method, packets sent from the IP address A to IP address B, from IP address A to IP address C, and from IP address C to IP address B could all use different ports in the channel.

Different load-balancing methods have different advantages, and the choice of a particular load-balancing method should be based on the position of the switch in the network and the kind of traffic that needs to be load-distributed. In [Figure 25-5](#), an EtherChannel of four workstations communicates with a router. Because the router is a single-MAC-address device, source-based forwarding on the switch EtherChannel ensures that the switch uses all available bandwidth to the router. The router is configured for destination-based forwarding because the large number of workstations ensures that the traffic is evenly distributed from the router 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 25-5 Load Distribution and Forwarding Methods



86482

EtherChannel and Switch Stacks

If a stack member that has ports participating in an EtherChannel fails or leaves the stack, the stack master removes the failed stack member switch ports from the EtherChannel. The remaining ports of the EtherChannel, if any, continue to provide connectivity.

When a switch is added to an existing stack, the new switch receives the running configuration from the stack master and updates itself with the EtherChannel-related stack configuration. The stack member also receives the operational information (the list of ports that are up and are members of a channel).

When two stacks merge that have EtherChannels configured between them, self-looped ports result. Spanning tree detects this condition and acts accordingly.

If the stack master fails or leaves the stack, a new stack master is elected. A spanning-tree reconvergence is not triggered unless there is a change in the EtherChannel bandwidth. The new stack master synchronizes the configuration of the stack members to that of the stack master.

For more information about switch stacks, see [Chapter 5, “Managing Switch Stacks.”](#)

Configuring EtherChannels

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

- [Default EtherChannel Configuration, page 25-9](#)
- [EtherChannel Configuration Guidelines, page 25-10](#)
- [Configuring Layer 2 EtherChannels, page 25-11](#) (required)
- [Configuring Layer 3 EtherChannels, page 25-13](#) (required)
- [Configuring EtherChannel Load Balancing, page 25-16](#) (optional)
- [Configuring the PAgP Learn Method and Priority, page 25-17](#) (optional)



Note

Make sure that the interfaces are correctly configured. For more information, see the “[EtherChannel Configuration Guidelines](#)” section on page 25-10.



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 25-2](#) shows the default EtherChannel configuration.

Table 25-2 Default EtherChannel Configuration

Feature	Default Setting
Channel groups	None assigned.
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.
Load balancing	Load distribution on the switch is based on the source-MAC address of the incoming packet.

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:

- More than 12 EtherChannels cannot be configured on a Catalyst 3750 switch stack.
- Configure an EtherChannel with up to eight Ethernet interfaces of the same type.
- 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
- Do not configure a Switched Port Analyzer (SPAN) destination as part of an EtherChannel.
- Do not configure a secure port as part of an EtherChannel.
- Do not configure a port that is an active member of an EtherChannel as an 802.1X port. If 802.1X is enabled 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.
- For cross-stack EtherChannel configurations, disable PAgP on all interfaces targeted for the EtherChannel by using the **channel-group *channel-group-number* mode on** interface configuration command. Before adding a stack member interface to an existing EtherChannel, manually disable PAgP on all the interfaces that are members of the channel group, and then manually configure the cross-stack EtherChannel.
- If cross-stack EtherChannel is configured and the switch stack partitions, loops and forwarding misbehaviors can occur.

Configuring Layer 2 EtherChannels

You configure Layer 2 EtherChannels by assigning interfaces to a channel group with the **channel-group** interface configuration command. This command automatically creates the port-channel logical interface.

If you enabled PAgP on an interface in the **auto** or **desirable** mode, you must reconfigure it for the **on** mode by using the **channel-group channel-group-number mode on** interface configuration command before adding this interface to a cross-stack EtherChannel. PAgP is not supported on cross-stack EtherChannels.

Beginning in privileged EXEC mode, follow these steps to assign a Layer 2 Ethernet interface to a Layer 2 EtherChannel. This procedure is required.

	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 }	<p>Assign the interface to a channel group, and specify the PAgP mode.</p> <p>For <i>channel-group-number</i>, the range is 1 to 12. 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"> • 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. You must use this keyword when EtherChannel members are from different switches in the switch stack (cross-stack EtherChannel). • non-silent—If your switch is connected to a partner that is PAgP-capable, configure the switch interface for nonsilent operation when the interface is in the auto or desirable mode. If you do not specify non-silent, 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. <p>For information on compatible modes for the switch and its partner, see the “PAgP Modes” section on page 25-5.</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 configure an EtherChannel on a single switch in the stack. It assigns Gigabit Ethernet interfaces 0/4 and 0/5 on stack member 2 as static-access ports in VLAN 10 to channel 5 with the PAgP mode **desirable**:

```
Switch# configure terminal
Switch(config)# interface range gigabitethernet2/0/4 -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 non-silent
Switch(config-if-range)# end
```

This example shows how to configure cross-stack EtherChannel. It assigns Gigabit Ethernet interfaces 0/4 and 0/5 on stack member 2 and Gigabit Ethernet interface 0/3 on stack member 3 as static-access ports in VLAN 10 to channel 5 with the PAgP mode disabled (**on**):

```
Switch# configure terminal
Switch(config)# interface range gigabitethernet2/0/4 -5
Switch(config-if-range)# switchport mode access
Switch(config-if-range)# switchport access vlan 10
Switch(config-if-range)# channel-group 5 mode on
Switch(config-if-range)# exit
Switch(config)# interface gigabitethernet3/0/3
Switch(config-if)# switchport mode access
Switch(config-if)# switchport access vlan 10
Switch(config-if)# channel-group 5 mode on
Switch(config-if)# exit
```

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 should first manually create the port-channel logical interface 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. This procedure is required.

	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 12.
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 25-14.

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. This procedure is required.

	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 }	<p>Assign the interface to a channel group, and specify the PAgP mode.</p> <p>For <i>channel-group-number</i>, the range is 1 to 12. 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 25-13.</p> <p>Each EtherChannel can consist of up to eight compatibly configured Ethernet interfaces.</p> <p>For mode, select one of these keywords:</p> <ul style="list-style-type: none"> • 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. You must use this keyword when EtherChannel members are from different switches in the switch stack (cross-stack EtherChannel). • non-silent—If your switch is connected to a partner that is PAgP capable, configure the switch interface for nonsilent operation when the interface is in the auto or desirable mode. If you do not specify non-silent, 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. <p>For information on compatible modes for the switch and its partner, see the “PAgP Modes” section on page 25-5.</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 configure an EtherChannel on a single switch in the stack. It assigns Gigabit Ethernet interfaces 0/4 and 0/5 on stack member 2 to channel 5 with the PAgP mode **desirable**:

```
Switch# configure terminal
Switch(config)# interface range gigabitethernet2/0/4 -5
Switch(config-if-range)# no ip address
Switch(config-if-range)# channel-group 5 mode desirable non-silent
Switch(config-if-range)# end
```

This example shows how to configure cross-stack EtherChannel. It assigns Gigabit Ethernet interfaces 0/4 and 0/5 on stack member 2 and Gigabit Ethernet interface 0/3 on stack member 3 to channel 5 with the PAgP mode disabled (**on**):

```
Switch# configure terminal
Switch(config)# interface range gigabitethernet2/0/4 -5
Switch(config-if-range)# no ip address
Switch(config-if-range)# channel-group 5 mode on
Switch(config-if-range)# exit
Switch(config)# interface gigabitethernet3/0/3
Switch(config-if)# no ip address
Switch(config-if)# channel-group 5 mode on
Switch(config-if)# exit
```

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 [“Load Balancing and Forwarding Methods”](#) section on page 25-6.

Beginning in privileged EXEC mode, follow these steps to configure EtherChannel load balancing. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	port-channel load-balance { dst-ip dst-mac src-dst-ip src-dst-mac src-ip src-mac }	Configure an EtherChannel load-balancing method. The default is src-mac . Select one of these keywords to determine the load-distribution method: <ul style="list-style-type: none"> • dst-ip—Load distribution is based on the destination-host IP address. • dst-mac—Load distribution is based on the destination-host MAC address of the incoming packet. • src-dst-ip—Load distribution is based on the source-and-destination host IP address. • src-dst-mac—Load distribution is based on the source and destination-host MAC address. • src-ip—Load distribution is based on the source-host IP address. • src-mac—Load distribution is based on the source-MAC address of the incoming packet.
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 (logical) ports. The learn method must be configured the same at both ends of the link.

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 when the local device is an aggregate-port learner. Therefore, you must manually set the learning method on the local device to learn addresses by physical ports. You also must set the load-distribution method to source-based distribution, so that any given source MAC address is always sent on the same physical port.

You also can 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 with the **pagp port-priority** interface configuration command. The higher the priority, the more likely that the port will be selected.



Note

The Catalyst 3750 switch 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 3750 switch is a physical learner (such as a Catalyst 1900 series switch), we recommend that you configure the Catalyst 3750 switch as a physical-port learner by using the **pagp learn-method physical-port** interface configuration command. Set the load-distribution method based on the source MAC address by using the **port-channel load-balance src-mac** global configuration command. The switch then sends packets to the Catalyst 1900 switch using the same interface in the EtherChannel from which it learned the source address. 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. This procedure is optional.

	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 25-16. 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.
Step 6	show running-config or show pagp <i>channel-group-number</i> 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.

Displaying EtherChannel and PAgP Status

To display EtherChannel and PAgP status information, use the privileged EXEC commands described in [Table 25-3](#):

Table 25-3 *Commands for Displaying EtherChannel and PAgP Status*

Command	Description
show etherchannel [<i>channel-group-number</i>] { brief detail load-balance port port-channel summary }	Displays EtherChannel information in a brief, 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.

1. You can clear PAgP channel-group information and traffic counters by using the **clear pagp** {*channel-group-number* [**counters**] | **counters**} privileged EXEC command.

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

