



Configuring EtherChannels

This chapter describes how to configure EtherChannels on Layer 2 interfaces on the Catalyst 2970 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.



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 24-1](#)
- [Configuring EtherChannels, page 24-7](#)
- [Displaying EtherChannel and PAgP Status, page 24-13](#)

Understanding EtherChannels


These sections describe how EtherChannels work:

- [EtherChannel Overview, page 24-2](#)
- [Port-Channel Interfaces, page 24-3](#)
- [Port Aggregation Protocol, page 24-3](#)
- [Load Balancing and Forwarding Methods, page 24-5](#)

EtherChannel Overview

An EtherChannel consists of individual Gigabit Ethernet links bundled into a single logical link as shown in [Figure 24-1](#).

Figure 24-1 Typical EtherChannel Configuration



The EtherChannel provides full-duplex bandwidth up to 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 Layer 2 interfaces. For Catalyst 2970 switches, the number of EtherChannels is limited to 12. For more information, see the [“EtherChannel Configuration Guidelines”](#) section on page 24-8.

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.


Port-Channel Interfaces

When you create a Layer 2 EtherChannel, a port-channel logical interface is involved. You can create the EtherChannel in these ways:

- Use the **channel-group** interface configuration command. This command automatically creates the port-channel logical interface when the channel group gets its first physical interface. The **channel-group** command binds the physical (10/100/1000 ports) and the logical ports together as shown in Figure 24-2.
- Use the **interface port-channel** *port-channel-number* global configuration command to manually create the port-channel logical interface. Then use the **channel-group** *channel-group-number* interface configuration 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.

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 24-2 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.

By using PAgP, the switch learns the identity of partners capable of supporting PAgP and the capabilities of each interface. It then dynamically groups similarly configured interfaces 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 24-1 shows the user-configurable EtherChannel PAgP modes for the **channel-group** interface configuration command.

Table 24-1 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.

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.

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 24-3](#), 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 24-3 Load Distribution and Forwarding Methods



900569

Configuring EtherChannels

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

- [Default EtherChannel Configuration, page 24-7](#)
- [EtherChannel Configuration Guidelines, page 24-8](#)
- [Configuring Layer 2 EtherChannels, page 24-8](#) (required)
- [Configuring EtherChannel Load Balancing, page 24-10](#) (optional)
- [Configuring the PAgP Learn Method and Priority, page 24-11](#) (optional)



Note

Make sure that the interfaces are correctly configured. For more information, see the [“EtherChannel Configuration Guidelines”](#) section on page 24-8.



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 24-2](#) shows the default EtherChannel configuration.

Table 24-2 *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 2970 switch.
- 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.

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.

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.
Step 4	channel-group <i>channel-group-number</i> mode { auto [non-silent] desirable [non-silent] on }	Assign the interface to a channel group, and specify the PAgP mode. For <i>channel-group-number</i> , the range is 1 to 12. Each EtherChannel can have of up to eight compatibly configured Ethernet interfaces. For mode , select one of these keywords: <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. • 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. For information on compatible modes for the switch and its partner, see the “PAgP Modes” section on page 24-4.
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 configure an EtherChannel. It assigns 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 non-silent
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 [“Load Balancing and Forwarding Methods” section on page 24-5](#).

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 2970 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 2970 switch is a physical learner (such as a Catalyst 1900 series switch), we recommend that you configure the Catalyst 2970 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 24-10. 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 24-3](#):

Table 24-3 *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.

