



CHAPTER 35

Configuring EtherChannels and Link-State Tracking

This chapter describes how to configure EtherChannels on Layer 2 and Layer 3 ports on the Cisco ME 3400E Ethernet Access 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 also describes how to configure link-state tracking.



Note

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

- [Understanding EtherChannels, page 35-1](#)
- [Configuring EtherChannels, page 35-9](#)
- [Displaying EtherChannel, PAgP, and LACP Status, page 35-22](#)
- [Understanding Link-State Tracking, page 35-22](#)
- [Configuring Link-State Tracking, page 35-24](#)
- [Displaying Link-State Tracking Status, page 35-25](#)


Understanding EtherChannels

- [EtherChannel Overview, page 35-2](#)
- [Port-Channel Interfaces, page 35-3](#)
- [Port Aggregation Protocol, page 35-4](#)
- [Link Aggregation Control Protocol, page 35-6](#)
- [EtherChannel On Mode, page 35-7](#)
- [Load Balancing and Forwarding Methods, page 35-7](#)

EtherChannel Overview

An EtherChannel consists of individual Fast Ethernet or Gigabit Ethernet links bundled into a single logical link as shown in [Figure 35-1](#).

Figure 35-1 Typical EtherChannel Configuration



The EtherChannel provides full-duplex bandwidth of up to 800 Mbps between your switch and another switch or host for Fast EtherChannel on a switch with 24 Fast Ethernet ports. For Gigabit EtherChannel, you can configure up to 8 Gbps (8 ports of 1 Gbps), depending on the number of supported Gigabit Ethernet interfaces.



Note

Only network node interfaces (NNIs) and enhanced network interfaces (ENIs) support Link Aggregation Control Protocol (LACP) or Port Aggregation Protocol (PAgP). Use the **port-type {eni | nni}** interface configuration command to configure a port as an ENI or NNI. The switch must be running the metro IP access image to allow configuring of more than four ports as NNIs.

Each EtherChannel can consist of up to eight compatibly configured Ethernet ports. All ports in each EtherChannel must be configured as either Layer 2 or Layer 3 ports. The number of EtherChannels is limited to 48. For more information, see the [“EtherChannel Configuration Guidelines” section on page 35-10](#). The EtherChannel Layer 3 ports 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 10, “Configuring Interfaces.”](#)

You can configure an EtherChannel in one of these modes: Port Aggregation Protocol (PAgP), Link Aggregation Control Protocol (LACP), or On mode. PAgP and LACP are available only on NNIs and ENIs. Configure both ends of the EtherChannel in the same mode:

- When you configure one end of an EtherChannel in either PAgP or LACP mode, the system negotiates with the other end of the channel to determine which ports should become active. Incompatible ports are suspended.
- When you configure an EtherChannel in the **on** mode, no negotiations take place. The switch forces all compatible ports to become active in the EtherChannel. The other end of the channel (on the other switch) must also be configured in the **on** mode; otherwise, packet loss can occur.

The local port is put into an independent state and continues to carry data traffic as would any other single link. The port configuration does not change, but the port does not participate in the EtherChannel.

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


- With Layer 2 ports, 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 ports, 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 ports, the **channel-group** command binds the physical port and the logical interface together as shown in [Figure 35-2](#).

Each EtherChannel has a port-channel logical interface numbered from 1 to 48. This port-channel interface number corresponds to the one specified with the **channel-group** interface configuration command.

Figure 35-2 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 ports assigned to the port-channel interface. Configuration changes applied to the physical port affect only the port to which you apply the configuration. To change the parameters of all ports in an EtherChannel, apply the configuration commands to the port-channel interface.

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 ports.

**Note**

PAgP is only available on network node interfaces (NNIs) and enhanced network interfaces (ENIs).

By using PAgP, the switch learns the identity of partners capable of supporting PAgP and the capabilities of each port. It then dynamically groups similarly configured ports into a single logical link (channel or aggregate port). Similarly configured ports are grouped based on hardware, administrative, and port parameter constraints. For example, PAgP groups the ports 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 35-1 shows the user-configurable EtherChannel PAgP modes for the **channel-group** interface configuration command on an NNI or ENI.

Table 35-1 *EtherChannel PAgP Modes*

Mode	Description
auto	Places a port into a passive negotiating state in which the port responds to PAgP packets it receives but does not start PAgP packet negotiation. This setting minimizes the transmission of PAgP packets.
desirable	Places a port into an active negotiating state in which the port starts negotiations with other ports by sending PAgP packets.

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

Both the **auto** and **desirable** modes enable ports to negotiate with partner ports to form an EtherChannel based on criteria such as port speed and, for Layer 2 EtherChannels, trunking state and VLAN numbers.

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

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

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

If your switch is connected to a partner that is PAgP-capable, you can configure the switch port 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 port to a channel group, and to use the port for transmission.

PAgP Interaction with Other Features

Cisco Discovery Protocol (CDP) sends and receives packets over the physical ports in the EtherChannel.



Note PAgP and CDP are only available on NNIs and ENIs. User network interfaces (UNIs) do not support PAgP or CDP.

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 ports that are up and have PAgP enabled for the auto or desirable mode.

Link Aggregation Control Protocol

The LACP is defined in IEEE 802.3ad standard and enables Cisco switches to manage Ethernet channels between switches that conform to the standard. LACP facilitates the automatic creation of EtherChannels by exchanging LACP packets between Ethernet ports.


Note

LACP is available only on NNIs and ENIs.

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

LACP Modes

Table 35-2 shows the user-configurable EtherChannel LACP modes for the **channel-group** interface configuration command on an NNI or ENI.

Table 35-2 EtherChannel LACP Modes

Mode	Description
active	Places a port into an active negotiating state in which the port starts negotiations with other ports by sending LACP packets.
passive	Places a port into a passive negotiating state in which the port responds to LACP packets that it receives, but does not start LACP packet negotiation. This setting minimizes the transmission of LACP packets.

Both the **active** and **passive LACP** modes enable ports to negotiate with partner ports to an EtherChannel based on criteria such as port speed and, for Layer 2 EtherChannels, trunking state and VLAN numbers.

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

- A port in the **active** mode can form an EtherChannel with another port that is in the **active** or **passive** mode.
- A port in the **passive** mode cannot form an EtherChannel with another port that is also in the **passive** mode because neither port starts LACP negotiation.

LACP Interaction with Other Features

The CDP sends and receives packets over the physical ports in the EtherChannel. Trunk ports send and receive LACP 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.

LACP sends and receives LACP PDUs only from ports that are up and have LACP enabled for the active or passive mode.

EtherChannel On Mode

EtherChannel **on** mode can be used to manually configure an EtherChannel. The **on** mode forces a port to join an EtherChannel without negotiations. It can be useful if the remote device does not support PAgP or LACP. With the **on** mode, a usable EtherChannel exists only when both ends of the link are configured in the **on** mode.



Note

For UNIs, the only available mode is **on**.

Ports that are configured in the **on** mode in the same channel group must have compatible port characteristics, such as speed and duplex. Ports that are not compatible are suspended, even though they are configured in the **on** mode.



Caution

You should use care when using the **on** mode. This is a manual configuration, and ports on both ends of the EtherChannel must have the same configuration. If the group is misconfigured, packet loss or spanning-tree loops can occur.

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 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.

On the ME-3400E switch, load distribution based on the destination host MAC address supports only four ports per EtherChannel. When you configure EtherChannel destination-MAC address load balancing, the traffic is balanced only among four ports in the channel group. If you configure more than four ports in an EtherChannel with destination host MAC address load distribution, only four of the ports receive distributed traffic. This limitation does not apply to the other load distribution methods.

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 35-3](#), 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 35-3 Load Distribution and Forwarding Methods

101239

Configuring EtherChannels

- [Default EtherChannel Configuration, page 35-10](#)
- [EtherChannel Configuration Guidelines, page 35-10](#)
- [Configuring Layer 2 EtherChannels, page 35-11](#) (required)
- [Configuring Layer 3 EtherChannels, page 35-14](#) (required)
- [Configuring EtherChannel Load Balancing, page 35-17](#) (optional)
- [Configuring the PAgP Learn Method and Priority, page 35-18](#) (optional)
- [Configuring LACP Hot-Standby Ports, page 35-19](#) (optional)

**Note**

Make sure that the ports are correctly configured. For more information, see the [“EtherChannel Configuration Guidelines”](#) section on page 35-10.

**Note**

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

Default EtherChannel Configuration

Table 35-3 shows the default EtherChannel configuration.

Table 35-3 *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 NNIs and ENIs.
PAgP priority	128 on all NNIs and ENIs.
LACP mode	No default.
LACP learn method	Aggregate-port learning on all NNIs and ENIs.
LACP port priority	32768 on all NNIs and ENIs.
LACP system priority	32768.
LACP system ID	LACP system priority and the switch MAC address.
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 ports are automatically disabled to avoid network loops and other problems. Follow these guidelines to avoid configuration problems:

- Do not try to configure more than 48 EtherChannels on the switch.
- Configure a PAgP EtherChannel including only NNIs or only ENIs.
- Configure a LACP EtherChannel including only NNIs or only ENIs.
- Configure all ports in an EtherChannel to operate at the same speeds and duplex modes.
- All ports in an EtherChannel must be the same type, either UNI, NNI, or ENI. You cannot mix port types in an EtherChannel.
- On UNIs, the EtherChannel mode must always be configured to **on**.
- Enable all ports in an EtherChannel. A port 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 ports in the EtherChannel. UNIs and ENIs are disabled by default. NNIs are enabled by default.
- 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

**Note**

Spanning Tree Protocol is only supported on NNIs or ENIs on which it has been specifically enabled.

- Do not configure a port to be a member of more than one EtherChannel group.
- Do not configure an EtherChannel in both the PAgP and LACP modes. EtherChannel groups running PAgP and LACP can coexist on the same switch. Individual EtherChannel groups can run either PAgP or LACP, but they cannot interoperate.

**Note**

PAgP and LACP are only available on NNIs and ENIs.

- If the switch is running the metro access image, you can have only four NNIs on the switch at the same time; therefore, only four ports in an EtherChannel can support LACP and PAgP at the same time. If the switch is running the metro IP access image, there is no limit to the number of NNIs that can be configured on the switch.
- Do not configure a Switched Port Analyzer (SPAN) destination port as part of an EtherChannel.
- Do not configure a secure port as part of an EtherChannel or the reverse.
- Do not configure a private-VLAN port as part of an EtherChannel.
- Do not configure a port that is an active or a not-yet-active member of an EtherChannel as an 802.1x port. If you try to enable 802.1x on an EtherChannel port, an error message appears, and 802.1x is not enabled.
- If EtherChannels are configured on switch interfaces, remove the EtherChannel configuration from the interfaces before globally enabling 802.1x on a switch by using the **dot1x system-auth-control** global configuration command.
- Do not enable link-state tracking on individual interfaces that will be part of a downstream Etherchannel interface. For Layer 2 EtherChannels:
 - Assign all ports in the EtherChannel to the same VLAN, or configure them as trunks. Ports with different native VLANs cannot form an EtherChannel.
 - If you configure an EtherChannel from trunk ports, verify that the trunking mode is the same on all the trunks. Inconsistent trunk modes on EtherChannel ports can have unexpected results.
 - An EtherChannel supports the same allowed range of VLANs on all the ports in a trunking Layer 2 EtherChannel. If the allowed range of VLANs is not the same, the ports do not form an EtherChannel even when PAgP is set to the **auto** or **desirable** mode.
 - NNIs or ENIs 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 ports 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 ports in the channel.

Configuring Layer 2 EtherChannels

You configure Layer 2 EtherChannels by assigning ports 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 port 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>	Specify a physical port, and enter interface configuration mode. Valid interfaces include physical ports. For a PAgP EtherChannel, you can configure up to eight ports of the same type and speed for the same group. For a LACP EtherChannel, you can configure up to 16 Ethernet ports of the same type. Up to eight ports can be active, and up to eight ports can be in standby mode. Note If the interface is a UNI, you must enter the port-type {eni nni} interface configuration command before configuring PAgP or LACP.
Step 3	no shutdown	Enable the port, if necessary. By default, UNIs and ENIs are disabled, and NNIs are enabled.
Step 4	switchport mode {access trunk} switchport access vlan <i>vlan-id</i>	Assign all ports as static-access ports in the same VLAN, or configure them as trunks. If you configure the port as a static-access port, assign it to only one VLAN. The range is 1 to 4094.

	Command	Purpose
Step 5	channel-group <i>channel-group-number</i> mode { auto [non-silent] desirable [non-silent] on } { active passive }	Assign the port to a channel group, and specify the PAgP or the LACP mode. For <i>channel-group-number</i>, the range is 1 to 48. Note For UNIs, the only available mode is on. For mode, select one of these keywords: • auto—Enables PAgP only if a PAgP device is detected. It places the port into a passive negotiating state, in which the port responds to PAgP packets it receives but does not start PAgP packet negotiation. • desirable—Unconditionally enables PAgP. It places the port into an active negotiating state, in which the port starts negotiations with other ports by sending PAgP packets. • on—Forces the port to channel without PAgP or LACP. With the on mode, a usable EtherChannel exists only when a port group in the on mode is connected to another port group in the on mode. • non-silent—(Optional) If your switch is connected to a partner that is PAgP-capable, configure the switch port for nonsilent operation when the port 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 port to a channel group, and to use the port for transmission. • active—Enables LACP only if a LACP device is detected. It places the port into an active negotiating state in which the port starts negotiations with other ports by sending LACP packets. • passive—Enables LACP on the port and places it into a passive negotiating state in which the port responds to LACP packets that it receives, but does not start LACP packet negotiation. For information on compatible modes for the switch and its partner, see the “PAgP Modes” section on page 35-5 and the “LACP Modes” section on page 35-6.
Step 6	end	Return to privileged EXEC mode.
Step 7	show running-config	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

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

This example shows how to configure an EtherChannel. It assigns two ports as static-access ports in VLAN 10 to channel 5 with the PAGP mode **desirable**:

```
Switch# configure terminal
Switch(config)# interface range fastethernet0/1 -2
Switch(config-if-range)# port-type nni
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 an EtherChannel. It assigns two ports as static-access ports in VLAN 10 to channel 5 with the LACP mode **active**:

```
Switch# configure terminal
Switch(config)# interface range gigabitethernet0/1 -2
Switch(config-if-range)# switchport mode access
Switch(config-if-range)# switchport access vlan 10
Switch(config-if-range)# channel-group 5 mode active
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 ports 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 port to an EtherChannel, you must delete the IP address from the physical port 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>	Specify the port-channel logical interface, and enter interface configuration mode. For <i>port-channel-number</i> , the range is 1 to 48.
Step 3	no switchport	Put the port-channel 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.

	Command	Purpose
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.
Step 8		Assign an Ethernet port to the Layer 3 EtherChannel. For more information, see the “Configuring the Physical Interfaces” section on page 35-15.

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 port 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>	Specify a physical port, and enter interface configuration mode. Valid interfaces include physical ports. For a PAgP EtherChannel, you can configure up to eight ports of the same type and speed for the same group. For a LACP EtherChannel, you can configure up to 16 Ethernet ports of the same type. Up to eight ports can be active, and up to eight ports can be in standby mode. Note If the interface is a UNI, you must enter the port-type {eni nni} interface configuration command before configuring PAgP or LACP.
Step 3	no shutdown	Enable the port, if necessary. By default, UNIs and ENIs are disabled, and NNIs are enabled.
Step 4	no ip address	Ensure that there is no IP address assigned to the physical port.
Step 5	no switchport	Put the port into Layer 3 mode.

Command	Purpose
Step 6 channel-group <i>channel-group-number</i> mode { auto [non-silent] desirable [non-silent] on } { active passive }	Assign the port to a channel group, and specify the PAgP or the LACP mode. For <i>channel-group-number</i>, the range is 1 to 48. 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 35-14. Note For UNIs, the only available mode is on. For mode, select one of these keywords: • auto—Enables PAgP only if a PAgP device is detected. It places the port into a passive negotiating state, in which the port responds to PAgP packets it receives but does not start PAgP packet negotiation. • desirable—Unconditionally enables PAgP. It places the port into an active negotiating state, in which the port starts negotiations with other ports by sending PAgP packets. • on—Forces the port to channel without PAgP or LACP. With the on mode, a usable EtherChannel exists only when a port group in the on mode is connected to another port group in the on mode. • non-silent—(Optional) If your switch is connected to a partner that is PAgP capable, configure the switch port for nonsilent operation when the port 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 port to a channel group, and to use the port for transmission. • active—Enables LACP only if a LACP device is detected. It places the port into an active negotiating state in which the port starts negotiations with other ports by sending LACP packets. • passive—Enables LACP on the port and places it into a passive negotiating state in which the port responds to LACP packets that it receives, but does not start LACP packet negotiation. For information on compatible modes for the switch and its partner, see the “PAgP Modes” section on page 35-5 and the “LACP Modes” section on page 35-6.
Step 7 end	Return to privileged EXEC mode.
Step 8 show running-config	Verify your entries.
Step 9 copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to configure an EtherChannel. It assigns two ports to channel 5 with the LACP mode **active**:

```
Switch# configure terminal
Switch(config)# interface range gigabitethernet0/1 -2
Switch(config-if-range)# no ip address
Switch(config-if-range)# no switchport
Switch(config-if-range)# channel-group 5 mode active
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 35-7](#).

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 load-distribution methods: • 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. Note When you enter the dst-mac keyword, the traffic is balanced only among four ports in the channel group. If you configure more than four ports in an EtherChannel with destination host MAC address load distribution, only four of the ports receive distributed traffic. This limitation does not apply to the other load distribution methods. • 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.



Note

PAgP is available only on NNIs and ENIs.

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 ports 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 port within the group for all transmissions and use other ports for hot standby. The unused ports in the group can be swapped into operation in just a few seconds if the selected single port loses hardware-signal detection. You can configure which port 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 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.

When the link partner to the switch is a physical learner, we recommend that you configure the Cisco ME 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 physical learner switch using the same port 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>	Specify the port for transmission, and enter interface configuration mode. Note If the interface is a UNI, you must enter the port-type {eni nni} interface configuration command before configuring LACP.
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 ports 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 35-17 . The learning method must be configured the same at both ends of the link. Note If the interface is a UNI, you must enter the port-type {eni nni} interface configuration command before configuring PAgP.
Step 4	pagp port-priority <i>priority</i>	Assign a priority so that the selected port 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 port 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.

Configuring LACP 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. The software places any additional links in a hot-standby mode. If one of the active links becomes inactive, a link that is in the hot-standby mode becomes active in its place.



Note

LACP is only available on NNIs and ENIs.

If you configure more than eight links for an EtherChannel group, the software automatically decides which of the hot-standby ports to make active based on the LACP priority. The software assigns to every link between systems that operate LACP a unique priority made up of these elements (in priority order):

- LACP system priority
- System ID (a combination of the LACP system priority and the switch MAC address)
- LACP port priority
- Port number

In priority comparisons, numerically lower values have higher priority. The priority decides which ports should be put in standby mode when there is a hardware limitation that prevents all compatible ports from aggregating.

Ports are considered for active use in aggregation in link-priority order starting with the port attached to the highest priority link. Each port is selected for active use if the preceding higher priority selections can also be maintained. Otherwise, the port is selected for standby mode.

You can change the default values of the LACP system priority and the LACP port priority to affect how the software selects active and standby links. For more information, see the [“Configuring the LACP System Priority” section on page 35-20](#) and the [“Configuring the LACP Port Priority” section on page 35-21](#).

Configuring the LACP System Priority

You can configure the system priority for all of the EtherChannels that are enabled for LACP by using the **lacp system-priority** global configuration command. You cannot configure a system priority for each LACP-configured channel. By changing this value from the default, you can affect how the software selects active and standby links.

You can use the **show etherchannel summary** privileged EXEC command to see which ports are in the hot-standby mode (denoted with an H port-state flag).

Beginning in privileged EXEC mode, follow these steps to configure the LACP system priority. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	lacp system-priority <i>priority</i>	Configure the LACP system priority. For <i>priority</i> , the range is 1 to 65535. The default is 32768. The lower the value, the higher the system priority.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config or show lacp sys-id	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the LACP system priority to the default value, use the **no lacp system-priority** global configuration command.

Configuring the LACP Port Priority

By default, all ports use the same port priority. If the local system has a lower value for the system priority and the system ID than the remote system, you can affect which of the hot-standby links become active first by changing the port priority of LACP EtherChannel ports to a lower value than the default. The hot-standby ports that have lower port numbers become active in the channel first. You can use the **show etherchannel summary** privileged EXEC command to see which ports are in the hot-standby mode (denoted with an H port-state flag).



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 the hot-standby state and are used only if one of the channeled ports fails.

Beginning in privileged EXEC mode, follow these steps to configure the LACP port priority. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the port to be configured, and enter interface configuration mode. Note If the interface is a UNI, you must enter the port-type {eni nni} interface configuration command before configuring LACP.
Step 3	lacp port-priority <i>priority</i>	Configure the LACP port priority. For <i>priority</i> , the range is 1 to 65535. The default is 32768. The lower the value, the more likely that the port will be used for LACP transmission.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config or show lacp [<i>channel-group-number</i>] internal	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the LACP port priority to the default value, use the **no lacp port-priority** interface configuration command.

Displaying EtherChannel, PAgP, and LACP Status

To display EtherChannel, PAgP, and LACP status information, use the privileged EXEC commands described in [Table 35-4](#):

Table 35-4 *Commands for Displaying EtherChannel, PAgP, and LACP Status*

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

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

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


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

Understanding Link-State Tracking

Link-state tracking, also known as trunk failover, is a feature that binds the link state of multiple interfaces. For example, link-state tracking provides redundancy in the network when used with Flex Links. If the link is lost on the primary interface, connectivity is transparently switched to the secondary interface.

As shown in [Figure 35-4](#), switches that could be Cisco ME3400E switches are used as user-facing provider edge (UPE) switches in a customer site at the edge of the provider network connected to a customer premises equipment (CPE) switch. The UPE switches are connected to the provider edge (PE) switches in the service provider (SP) network. Customer devices, such as clients, connected to the CPE switch have multiple connections to the SP network. This configuration ensures that the traffic flow is balanced from the customer site to the SP and the reverse. Ports connected to the CPE are referred to as downstream ports, and ports connected to PE switches are referred to as upstream ports.

- UPE switch A provides links to the CPE through link-state group 1. Port 1 and port 2 are connected to the CPE. Port 3 and port 4 are connected to PE switch A through link-state group 1.
- UPE switch B provides links to the CPE through link-state group 2. Port 1 and port 2 are connected to CPE. Port 3 and 4 are connected to PE switch A through link-state group 2.

Figure 35-4 Typical Link-State Tracking Configuration

141683

When you enable link-state tracking on the switch, the link state of the downstream ports is bound to the link state of one or more of the upstream ports. After you associate a set of downstream ports to a set of upstream ports, if all of the upstream ports become unavailable, link-state tracking automatically puts the associated downstream ports in an error-disabled state. This causes the CPE primary interface to failover to the secondary interface.

If the PE switch fails, the cables are disconnected, or the link is lost, the upstream interfaces can lose connectivity. When link-state tracking is not enabled and the upstream interfaces lose connectivity, the link states of the downstream interfaces remain unchanged. The CPE is not aware that upstream connectivity has been lost and does not failover to the secondary interface.

An interface can be an aggregation of ports (an EtherChannel), a single physical port in access or trunk mode, or routed ports. These interfaces can be bundled together, and each downstream interface can be associated with a single group consisting of multiple upstream interfaces, referred to as a link-state group.

In a link-state group, the link state of the downstream interfaces is dependent on the link state of the upstream interfaces. If all of the upstream interfaces in a link-state group are in the link-down state, the associated downstream interfaces are forced into the link-down state. If any one of the upstream interfaces in the link-state group is in the link-up state, the associated downstream interfaces can change to or remain in a link-up state.

For example, in [Figure 35-4](#), downstream interfaces 1 and 2 on UPE switch A are defined in link-state group 1 with upstream interfaces 3 and 4. Similarly, downstream interfaces 1 and 2 on UPE switch B are defined in link-state group 2 with upstream interfaces 3 and 4.

If the link is lost on upstream interface 3, the link states of downstream interfaces 1 and 2 do not change. If upstream interface 4 also loses link, downstream interfaces 1 and 2 change to the link-down state. The CPE switch stops forwarding traffic to PE switch A and starts to forward traffic to PE switch B.

You can recover a downstream interface link-down condition by removing the failed downstream port from the link-state group. To recover multiple downstream interfaces, disable the link-state group.

Configuring Link-State Tracking

- [Default Link-State Tracking Configuration, page 35-24](#)
- [Link-State Tracking Configuration Guidelines, page 35-24](#)
- [Configuring Link-State Tracking, page 35-24](#)

Default Link-State Tracking Configuration

There are no link-state groups defined, and link-state tracking is not enabled for any group.

Link-State Tracking Configuration Guidelines

- An interface that is defined as an upstream interface cannot also be defined as a downstream interface in the same or a different link-state group. The reverse is also true.
- An interface cannot be a member of more than one link-state group.
- You can configure only two link-state groups per switch.
- Do not enable link-state tracking on individual interfaces that will be part of a downstream Etherchannel interface.

Configuring Link-State Tracking

Beginning in privileged EXEC mode, follow these steps to configure a link-state group and to assign an interface to a group:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	link state track <i>number</i>	Create a link-state group, and enable link-state tracking. The group number can be 1 to 2; the default is 1.
Step 3	interface <i>interface-id</i>	Specify a physical interface or range of interfaces to configure, and enter interface configuration mode. Valid interfaces include switch ports in access or trunk mode (IEEE 802.1q), routed ports, or multiple ports bundled into an upstream EtherChannel interface (static, PAgP, or LACP), also in trunk mode. Note Do not enable link-state tracking on individual interfaces that will be part of a downstream Etherchannel interface.
Step 4	no shutdown	Enable the port, if necessary. By default, UNIs and ENIs are disabled, and NNIs are enabled.

	Command	Purpose
Step 5	link state group [<i>number</i>] { upstream downstream }	Specify a link-state group, and configure the interface as either an upstream or downstream interface in the group. The group number can be 1 to 2; the default is 1.
Step 6	end	Return to privileged EXEC mode.
Step 7	show running-config	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to create a link-state group and configure the interfaces:

```
Switch# configure terminal
Switch(config)# link state track 1
Switch(config)# interface range fastethernet/0/9 -10
Switch(config-if)# link state group 1 upstream
Switch(config-if)# interface fastethernet0/1
Switch(config-if)# link state group 1 downstream
Switch(config-if)# interface fastethernet0/3
Switch(config-if)# link state group 1 downstream
Switch(config-if)# interface fastethernet0/5
Switch(config-if)# link state group 1 downstream
Switch(config-if)# end
```

To disable a link-state group, use the **no link state track** *number* global configuration command.

Displaying Link-State Tracking Status

Use the **show link state group** command to display the link-state group information. Enter this command without keywords to display information about all link-state groups. Enter the group number to display information specific to the group. Enter the detail keyword to display detailed information about the group.

This is an example of output from the **show link state group 1** command:

```
Switch> show link state group 1

Link State Group: 1      Status: Enabled, Down
```

This is an example of output from the **show link state group detail** command:

```
Switch> show link state group detail

(Up):Interface up      (Dwn):Interface Down    (Dis):Interface disabled

Link State Group: 1 Status: Enabled, Down
Upstream Interfaces : Fa0/15(Dwn) Fa0/16(Dwn)
Downstream Interfaces : Fa0/11(Dis) Fa0/12(Dis) Fa0/13(Dis) Fa0/14(Dis)

Link State Group: 2 Status: Enabled, Down
Upstream Interfaces : Fa0/15(Dwn) Fa0/16(Dwn) Fa0/17(Dwn)
Downstream Interfaces : Fa0/11(Dis) Fa0/12(Dis) Fa0/13(Dis) Fa0/14(Dis)

(Up):Interface up (Dwn):Interface Down (Dis):Interface disabled
```

For detailed information about the fields in the display, see the command reference for this release.

