



Configuring Fallback Bridging

This chapter describes how to configure fallback bridging (VLAN bridging) on your Catalyst 3550 switch. With fallback bridging, you can forward non-IP packets that the switch does not route between VLAN bridge domains and routed ports.

To use this feature, you must have the enhanced multilayer (EMI) image installed on your switch.



Note

For complete syntax and usage information for the commands used in this chapter, see the *Cisco IOS Bridging and IBM Networking Command Reference, Volume 1 of 2, Release 12.2*.

This chapter consists of these sections:

- [Understanding Fallback Bridging, page 36-1](#)
- [Configuring Fallback Bridging, page 36-3](#)
- [Monitoring and Maintaining Fallback Bridging, page 36-12](#)

Understanding Fallback Bridging

With fallback bridging, the switch bridges together two or more VLANs or routed ports, essentially connecting multiple VLANs within one bridge domain. Fallback bridging forwards traffic that the switch does not route and forwards traffic belonging to a nonroutable protocol such as DECnet.

Fallback bridging does not allow the spanning trees from the VLANs being bridged to collapse; each VLAN has its own spanning-tree instance and a separate spanning tree, called the VLAN-bridge spanning tree, which runs on top of the bridge group to prevent loops.

A VLAN bridge domain is represented with switch virtual interface (SVI). A set of SVIs and routed ports (which do not have any VLANs associated with them) can be configured (grouped together) to form a bridge group. Recall that an SVI represents a VLAN of switch ports as one interface to the routing or bridging function in the system. You associate only one SVI with a VLAN, and you configure an SVI for a VLAN only when you want to route between VLANs, to fallback-bridge nonroutable protocols between VLANs, or to provide IP host connectivity to the switch. A routed port is a physical port that acts like a port on a router, but it is not connected to a router. A routed port is not associated with a particular VLAN, does not support VLAN subinterfaces, but behaves like a normal routed interface. For more information about SVIs and routed ports, see [Chapter 9, “Configuring Interface Characteristics.”](#)


A bridge group is an internal organization of network interfaces on a switch. Bridge groups cannot be used to identify traffic switched within the bridge group outside the switch on which they are defined. Bridge groups on the same switch function as distinct bridges; that is, bridged traffic and bridge protocol data units (BPDUs) are not exchanged between different bridge groups on a switch. An interface can be a member of only one bridge group. Use a bridge group for each separately bridged (topologically distinct) network connected to the switch.

These are the reasons for placing network interfaces into a bridge group:

- To bridge all nonrouted traffic among the network interfaces making up the bridge group. If the packet destination address is in the bridge table, the packet is forwarded on a single interface in the bridge group. If the packet destination address is not in the bridge table, the packet is flooded on all forwarding interfaces in the bridge group. The switch places source addresses in the bridge table as it learns them during the bridging process.
- To participate in the spanning-tree algorithm by receiving, and in some cases sending, BPDUs on the LANs to which they are attached. A separate spanning-tree process runs for each configured bridge group. Each bridge group participates in a separate spanning-tree instance. A bridge group establishes a spanning-tree instance based on the BPDUs it receives on only its member interfaces.

Figure 36-1 shows a fallback bridging network example. The switch has two interfaces configured as SVIs with different assigned IP addresses and attached to two different VLANs. Another interface is configured as a routed port with its own IP address. If all three of these ports are assigned to the same bridge group, non-IP protocol frames can be forwarded among the end stations connected to the switch even though they are on different networks and in different VLANs. IP addresses do not need to be assigned to routed ports or SVIs for fallback bridging to work.

Figure 36-1 Fallback Bridging Network Example



Configuring Fallback Bridging

These sections describe how to configure fallback bridging on your switch:

- [Default Fallback Bridging Configuration, page 36-3](#)
- [Fallback Bridging Configuration Guidelines, page 36-3](#)
- [Creating a Bridge Group, page 36-4](#) (required)
- [Preventing the Forwarding of Dynamically Learned Stations, page 36-5](#) (optional)
- [Configuring the Bridge Table Aging Time, page 36-6](#) (optional)
- [Filtering Frames by a Specific MAC Address, page 36-6](#) (optional)
- [Adjusting Spanning-Tree Parameters, page 36-7](#) (optional)

Default Fallback Bridging Configuration

[Table 36-1](#) shows the default fallback bridging configuration.

Table 36-1 *Default Fallback Bridging Configuration*

Feature	Default Setting
Bridge groups	None are defined or assigned to an interface. No VLAN-bridge STP is defined.
Switch forwards frames for stations that it has dynamically learned	Enabled.
Bridge table aging time for dynamic entries	300 seconds.
MAC-layer frame filtering	Disabled.
Spanning tree parameters:	
<ul style="list-style-type: none"> • Switch priority • Interface priority • Interface path cost 	<ul style="list-style-type: none"> • 32768. • 128. • 10 Mbps: 100. 100 Mbps: 19. 1000 Mbps: 4.
<ul style="list-style-type: none"> • Hello BPDU interval • Forward-delay interval • Maximum idle interval 	<ul style="list-style-type: none"> • 2 seconds. • 20 seconds. • 30 seconds.

Fallback Bridging Configuration Guidelines

A maximum of 31 bridge groups can be configured on the switch.

An interface (an SVI or routed port) can be a member of only one bridge group.

Use a bridge group for each separately bridged (topologically distinct) network connected to the switch.

All protocols except IP (Version 4 and Version 6), Address Resolution Protocol (ARP), reverse ARP (RARP), LOOPBACK, and Frame Relay ARP are fallback bridged.

Creating a Bridge Group

To configure fallback bridging for a set of SVIs or routed ports, these interfaces must be assigned to bridge groups. All interfaces in the same group belong to the same bridge domain. Each SVI or routed port can be assigned to only one bridge group. A maximum of 31 bridge groups can be configured on the switch.



Note

The protected port feature is not compatible with fallback bridging. When fallback bridging is enabled, it is possible for packets to be forwarded from one protected port on a switch to another protected port on the same switch if the ports are in different VLANs.

Beginning in privileged EXEC mode, follow these steps to create a bridge group and assign an interface to it. This procedure is required.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	bridge <i>bridge-group</i> protocol vlan-bridge	Assign a bridge group number, and specify the VLAN-bridge spanning-tree protocol to run in the bridge group. The ibm and dec keywords are not supported. For <i>bridge-group</i> , specify the bridge group number. The range is 1 to 255. You can create up to 31 bridge groups. Frames are bridged only among interfaces in the same group.
Step 3	interface <i>interface-id</i>	Specify the interface on which you want to assign the bridge group, and enter interface configuration mode. The specified interface must be one of these: <ul style="list-style-type: none"> • A routed port: a physical port that you have configured as a Layer 3 port by entering the no switchport interface configuration command. • An SVI: a VLAN interface that you created by using the interface vlan <i>vlan-id</i> global configuration command. Note You can assign an IP address to the routed port or to the SVI, but it is not required.
Step 4	bridge-group <i>bridge-group</i>	Assign the interface to the bridge group created in Step 2. By default, the interface is not assigned to any bridge group. An interface can be assigned to only one bridge group.
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 a bridge group, use the **no bridge** *bridge-group* global configuration command. To remove an interface from a bridge group and to remove the bridge group, use the **no bridge-group** *bridge-group* interface configuration command.

This example shows how to create bridge group 10, to specify that the VLAN-bridge STP runs in the bridge group, to define an interface as a routed port, and to assign the interface to the bridge group:

```
Switch(config)# bridge 10 protocol vlan-bridge
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# no switchport
Switch(config-if)# no shutdown
Switch(config-if)# bridge-group 10
```

This example shows how to create bridge group 10 and to specify that the VLAN-bridge STP runs in the bridge group. It defines an interface as an SVI and assigns this interface to VLAN 2 and to the bridge group:

```
Switch(config)# bridge 10 protocol vlan-bridge
Switch(config)# vlan 2
Switch(config-vlan)# exit
Switch(config)# interface vlan2
Switch(config-if)# exit
Switch(config)# interface gigabitethernet0/2
Switch(config-if)# switchport mode access
Switch(config-if)# switchport access vlan 2
Switch(config-if)# no switchport
Switch(config-if)# no shutdown
Switch(config-if)# bridge-group 10
```

Preventing the Forwarding of Dynamically Learned Stations

By default, the switch forwards any frames for stations that it has dynamically learned. By disabling this activity, the switch only forwards frames whose addresses have been statically configured into the forwarding cache.

Beginning in privileged EXEC mode, follow these steps to prevent the switch from forwarding frames for stations that it has dynamically learned. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no bridge <i>bridge-group</i> acquire	Enable the switch to stop forwarding any frames for stations that it has dynamically learned through the discovery process and to limit frame forwarding to statically configured stations. The switch filters all frames except those whose destined-to addresses have been statically configured into the forwarding cache. To configure a static address, use the bridge <i>bridge-group</i> address <i>mac-address</i> {forward discard} global configuration command. For <i>bridge-group</i> , specify the bridge group number. The range is 1 to 255.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entry.
Step 5	copy running-config startup-config	(Optional) Save your entry in the configuration file.

To cause the switch to forward frames to stations that it has dynamically learned, use the **bridge *bridge-group* acquire** global configuration command.

This example shows how to prevent the switch from forwarding frames for stations that it has dynamically learned in bridge group 10:

```
Switch(config)# no bridge 10 acquire
```

Configuring the Bridge Table Aging Time

A switch forwards, floods, or drops packets based on the bridge table. The bridge table maintains both static and dynamic entries. Static entries are entered by you or learned by the switch. Dynamic entries are entered by the bridge learning process. A dynamic entry is automatically removed after a specified length of time, known as aging time, from the time the entry was created or last updated.

If you are likely to move hosts on a switched network, decrease the aging-time to enable the switch to quickly adapt to the change. If hosts on a switched network do not continuously send packets, increase the aging time to keep the dynamic entries for a longer time and thus reduce the possibility of flooding when the hosts send again.

Beginning in privileged EXEC mode, follow these steps to configure the aging time. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	bridge <i>bridge-group</i> aging-time <i>seconds</i>	Specify the length of time that a dynamic entry remains in the bridge table from the time the entry was created or last updated. <ul style="list-style-type: none"> For <i>bridge-group</i>, specify the bridge group number. The range is 1 to 255. For <i>seconds</i>, enter a number from 0 to 1000000. The default is 300 seconds.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entry.
Step 5	copy running-config startup-config	(Optional) Save your entry in the configuration file.

To return to the default aging-time interval, use the **no bridge *bridge-group* aging-time** global configuration command.

This example shows how to change the bridge table aging time to 200 seconds for bridge group 10:

```
Switch(config)# bridge 10 aging-time 200
```

Filtering Frames by a Specific MAC Address

A switch examines frames and sends them through the internetwork according to the destination address; a switch does not forward a frame back to its originating network segment. You can use the software to configure specific administrative filters that filter frames based on information other than the paths to their destinations.

You can filter frames with a particular MAC-layer station destination address. You can configure any number of addresses in the system without a performance penalty.

Beginning in privileged EXEC mode, follow these steps to filter by the MAC-layer address. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	bridge <i>bridge-group</i> address <i>mac-address</i> { forward discard } [<i>interface-id</i>]	Specify the MAC address to discard or forward. <ul style="list-style-type: none"> For <i>bridge-group</i>, specify the bridge group number. The range is 1 to 255. For address <i>mac-address</i>, specify the MAC-layer destination address to be filtered. Specify forward if you want the frame destined to the specified interface to be forwarded. Specify discard if you want the frame to be discarded. (Optional) For <i>interface-id</i>, specify the interface on which the address can be reached.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entry.
Step 5	copy running-config startup-config	(Optional) Save your entry in the configuration file.

To disable the frame forwarding ability, use the **no bridge** *bridge-group* **address** *mac-address* global configuration command.

This example shows how to forward a frame with MAC address 0800.cb00.45e9 through an interface in bridge group 1:

```
Switch(config)# bridge 1 address 0800.cb00.45e9 forward gigabitethernet0/1
```

Adjusting Spanning-Tree Parameters

You might need to adjust certain spanning-tree parameters if the default values are not suitable. You configure parameters affecting the entire spanning tree by using variations of the **bridge** global configuration command. You configure interface-specific parameters by using variations of the **bridge-group** interface configuration command.

You can adjust spanning-tree parameters by performing any of the tasks in these sections:

- [Changing the Switch Priority, page 36-8](#) (optional)
- [Changing the Interface Priority, page 36-8](#) (optional)
- [Assigning a Path Cost, page 36-9](#) (optional)
- [Adjusting BPDU Intervals, page 36-10](#) (optional)
- [Disabling the Spanning Tree on an Interface, page 36-12](#) (optional)



Note

Only network administrators with a good understanding of how switches and STP function should make adjustments to spanning-tree parameters. Poorly planned adjustments can have a negative impact on performance. A good source on switching is the IEEE 802.1D specification; for more information, see the “References and Recommended Reading” appendix in the *Cisco IOS Configuration Fundamentals Command Reference*.

Changing the Switch Priority

You can globally configure the priority of an individual switch when two switches tie for position as the root switch, or you can configure the likelihood that a switch will be selected as the root switch. This priority is determined by default; however, you can change it.

Beginning in privileged EXEC mode, follow these steps to change the switch priority. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	bridge <i>bridge-group</i> priority <i>number</i>	Change the priority of the switch. <ul style="list-style-type: none"> For <i>bridge-group</i>, specify the bridge group number. The range is 1 to 255. For <i>number</i>, enter a number from 0 to 65535. The default is 32768. The lower the number, the more likely the switch will be chosen as the root.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entry.
Step 5	copy running-config startup-config	(Optional) Save your entry in the configuration file.

To return to the default setting, use the **no bridge** *bridge-group* **priority** global configuration command. To change the priority on an interface, use the **bridge-group** **priority** interface configuration command (described in the next section).

This example shows how to set the switch priority to 100 for bridge group 10:

```
Switch(config)# bridge 10 priority 100
```

Changing the Interface Priority

You can change the priority for an interface. When two switches tie for position as the root switch, you configure an interface priority to break the tie. The switch with the lowest interface value is elected.

Beginning in privileged EXEC mode, follow these steps to change the interface priority. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface to set the priority, and enter interface configuration mode.
Step 3	bridge-group <i>bridge-group</i> priority <i>number</i>	Change the priority of an interface. <ul style="list-style-type: none"> For <i>bridge-group</i>, specify the bridge group number. The range is 1 to 255. For <i>number</i>, enter a number from 0 to 255. The lower the number, the more likely that the interface on the switch will be chosen as the root. The default is 128.
Step 4	end	Return to privileged EXEC mode.

	Command	Purpose
Step 5	<code>show running-config</code>	Verify your entry.
Step 6	<code>copy running-config startup-config</code>	(Optional) Save your entry in the configuration file.

No **no** form of this command exists. To return to the default setting, use the **no bridge-group *bridge-group* priority** interface configuration command.

This example shows how to change the priority of an interface to 20 in bridge group 10:

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# bridge-group 10 priority 20
```

Assigning a Path Cost

Each interface has a path cost associated with it. By convention, the path cost is 1000/data rate of the attached LAN, in Mbps.

Beginning in privileged EXEC mode, follow these steps to assign a path cost. This procedure is optional.

	Command	Purpose
Step 1	<code>configure terminal</code>	Enter global configuration mode.
Step 2	<code>interface <i>interface-id</i></code>	Specify the interface to set the path cost, and enter interface configuration mode.
Step 3	<code>bridge-group <i>bridge-group</i> path-cost <i>cost</i></code>	Assign the path cost of an interface. <ul style="list-style-type: none"> For <i>bridge-group</i>, specify the bridge group number. The range is 1 to 255. For <i>cost</i>, enter a number from 1 to 65536. The higher the value, the higher the cost. <ul style="list-style-type: none"> For 10 Mbps, the default path cost is 100. For 100 Mbps, the default path cost is 19. For 1000 Mbps, the default path cost is 4.
Step 4	<code>end</code>	Return to privileged EXEC mode.
Step 5	<code>show running-config</code>	Verify your entry.
Step 6	<code>copy running-config startup-config</code>	(Optional) Save your entry in the configuration file.

To return to the default path cost, use the **no bridge-group *bridge-group* path-cost** interface configuration command.

This example shows how to change the path cost on an interface to 20 in bridge group 10:

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# bridge-group 10 path-cost 20
```

Adjusting BPDU Intervals

You can adjust BPDU intervals as described in these sections:

- [Adjusting the Interval between Hello BPDUs, page 36-10](#) (optional)
- [Changing the Forward-Delay Interval, page 36-10](#) (optional)
- [Changing the Maximum-Idle Interval, page 36-11](#) (optional)



Note

Each switch in a spanning tree adopts the interval between hello BPDUs, the forward delay interval, and the maximum idle interval parameters of the root switch, regardless of what its individual configuration might be.

Adjusting the Interval between Hello BPDUs

Beginning in privileged EXEC mode, follow these step to adjust the interval between hello BPDUs. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	bridge <i>bridge-group</i> hello-time <i>seconds</i>	Specify the interval between hello BPDUs. <ul style="list-style-type: none"> • For <i>bridge-group</i>, specify the bridge group number. The range is 1 to 255. • For <i>seconds</i>, enter a number from 1 to 10. The default is 2 seconds.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entry.
Step 5	copy running-config startup-config	(Optional) Save your entry in the configuration file.

To return to the default setting, use the **no bridge** *bridge-group* **hello-time** global configuration command.

This example shows how to change the hello interval to 5 seconds in bridge group 10:

```
Switch(config)# bridge 10 hello-time 5
```

Changing the Forward-Delay Interval

The forward-delay interval is the amount of time spent listening for topology change information after an interface has been activated for switching and before forwarding actually begins.

Beginning in privileged EXEC mode, follow these steps to change the forward-delay interval. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	bridge <i>bridge-group</i> forward-time <i>seconds</i>	Specify the forward-delay interval. <ul style="list-style-type: none"> For <i>bridge-group</i>, specify the bridge group number. The range is 1 to 255. For <i>seconds</i>, enter a number from 10 to 200. The default is 20 seconds.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entry.
Step 5	copy running-config startup-config	(Optional) Save your entry in the configuration file.

To return to the default setting, use the **no bridge *bridge-group* forward-time** global configuration command.

This example shows how to change the forward-delay interval to 10 seconds in bridge group 10:

```
Switch(config)# bridge 10 forward-time 10
```

Changing the Maximum-Idle Interval

If a switch does not receive BPDUs from the root switch within a specified interval, it recomputes the spanning-tree topology.

Beginning in privileged EXEC mode, follow these steps to change the maximum-idle interval (maximum aging time). This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	bridge <i>bridge-group</i> max-age <i>seconds</i>	Specify the interval that the switch waits to hear BPDUs from the root switch. <ul style="list-style-type: none"> For <i>bridge-group</i>, specify the bridge group number. The range is 1 to 255. For <i>seconds</i>, enter a number from 10 to 200. The default is 30 seconds.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entry.
Step 5	copy running-config startup-config	(Optional) Save your entry in the configuration file.

To return to the default setting, use the **no bridge *bridge-group* max-age** global configuration command.

This example shows how to change the maximum-idle interval to 30 seconds in bridge group 10:

```
Switch(config)# bridge 10 max-age 30
```

Disabling the Spanning Tree on an Interface

When a loop-free path exists between any two switched subnetworks, you can prevent BPDUs generated in one switching subnetwork from impacting devices in the other switching subnetwork, yet still permit switching throughout the network as a whole. For example, when switched LAN subnetworks are separated by a WAN, BPDUs can be prevented from traveling across the WAN link.

Beginning in privileged EXEC mode, follow these steps to disable spanning tree on an interface. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface, and enter interface configuration mode.
Step 3	bridge-group <i>bridge-group</i> spanning-disabled	Disable spanning tree on the interface. For <i>bridge-group</i> , specify the bridge group number. The range is 1 to 255.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entry.
Step 6	copy running-config startup-config	(Optional) Save your entry in the configuration file.

To re-enable spanning tree on the interface, use the **no bridge-group** *bridge-group* **spanning-disabled** interface configuration command.

This example shows how to disable spanning tree on an interface in bridge group 10:

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# bridge group 10 spanning-disabled
```

Monitoring and Maintaining Fallback Bridging

To monitor and maintain fallback bridging, use one or more of the privileged EXEC commands in [Table 36-2](#):

Table 36-2 Commands for Monitoring and Maintaining Fallback Bridging

Command	Purpose
clear bridge <i>bridge-group</i>	Removes any learned entries from the forwarding database and clears the transmit and receive counts for any statically configured entries.
show bridge [<i>bridge-group</i>]	Displays details about the bridge group.
show bridge [<i>bridge-group</i>] [<i>interface-id</i>] [<i>address</i>] [group] [verbose]	Displays classes of entries in the bridge forwarding database.

For information about the fields in these displays, see the *Cisco IOS Bridging and IBM Networking Command Reference, Volume 1 of 2, Release 12.2*.