CHAPTER 54

Configuring Fallback Bridging

This chapter describes how to configure fallback bridging (VLAN bridging) on the Catalyst 3750-X or 3560-X 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, the switch or stack master must be running the IP services feature set. Unless otherwise noted, the term *switch* refers to a Catalyst 3750-X or 3560-X standalone switch and to a Catalyst 3750-X switch stack.

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 12.4*.

- Understanding Fallback Bridging, page 54-1
- Configuring Fallback Bridging, page 54-3
- Monitoring and Maintaining Fallback Bridging, page 54-10

Understanding Fallback Bridging

- Fallback Bridging Overview, page 54-1
- Fallback Bridging and Switch Stacks, page 54-3

Fallback Bridging Overview

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.

A VLAN bridge domain is represented with switch virtual interfaces (SVIs). 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 port. For more information about SVIs and routed ports, see Chapter 15, “Configuring Interface Characteristics.”

A bridge group is an internal organization of network interfaces on a switch. You cannot use bridge groups to identify traffic switched within the bridge group outside the switch on which they are defined. Bridge groups on the 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.

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.

The switch creates a VLAN-bridge spanning-tree instance when a bridge group is created. The switch runs the bridge group and treats the SVIs and routed ports in the bridge group as its spanning-tree ports.

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. A source MAC address is learned on a bridge group only when the address is learned on a VLAN (the reverse is not true). Any address that is learned on a stack member is learned by all switches in the stack.
- 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. If the bridge STP BPDU is received on a port whose VLAN does not belong to a bridge group, the BPDU is flooded on all the forwarding ports of the VLAN.

Figure 54-1 shows a fallback bridging network example. The switch has two ports configured as SVIs with different assigned IP addresses and attached to two different VLANs. Another port 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 54-1 Fallback Bridging Network Example

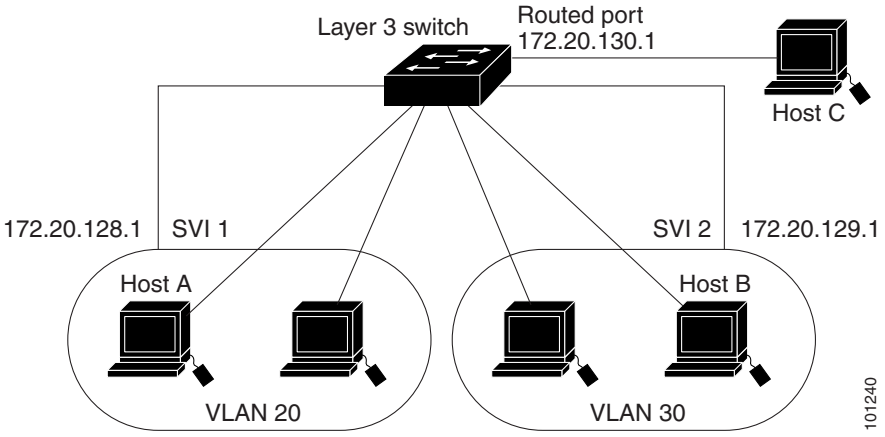

Fallback Bridging and Switch Stacks

When the stack master fails, a stack member becomes the new stack master by using the election process described in Chapter 5, "Managing Switch Stacks." The new stack master creates new VLAN-bridge spanning-tree instance, which temporarily puts the spanning-tree ports used for fallback bridging into a nonforwarding state. A momentary traffic disruption occurs until the spanning-tree states transition to the forwarding state. All MAC addresses must be relearned in the bridge group.

Note If a stack master running the IP services feature set fails and if the newly elected stack master is running the IP base feature set, the switch stack loses its fallback bridging capability.

If stacks merge or if a switch is added to the stack, any new VLANs that are part of a bridge group and become active are included in the VLAN-bridge STP.

When a stack member fails, the addresses learned from this member are deleted from the bridge group MAC address table.

For more information about switch stacks, see Chapter 5, "Managing Switch Stacks."

Configuring Fallback Bridging

- Default Fallback Bridging Configuration, page 54-3
- Fallback Bridging Configuration Guidelines, page 54-4
- Creating a Bridge Group, page 54-4 (required)
- Adjusting Spanning-Tree Parameters, page 54-5 (optional)

Default Fallback Bridging Configuration

Table 54-1 Default Fallback Bridging Configuration

Feature	Default Setting
Bridge groups	None are defined or assigned to a port. No VLAN-bridge STP is defined.
Switch forwards frames for stations that it has dynamically learned	Enabled.
Spanning tree parameters:	
• Switch priority	• 32768.
• Port priority	• 128.
• Port path cost	• 10 Mb/s: 100. 100 Mb/s: 19. 1000 Mb/s: 4.
• Hello BPDU interval	• 2 seconds.
• Forward-delay interval	• 20 seconds.
• Maximum idle interval	• 30 seconds.

Fallback Bridging Configuration Guidelines

Up to 32 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.

Do not configure fallback bridging on a switch configured with private VLANs.

All protocols except IP (Version 4 and Version 6), Address Resolution Protocol (ARP), reverse ARP (RARP), LOOPBACK, Frame Relay ARP, and shared STP packets 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.

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 to assign an interface to it. This procedure is required.

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**bridge** *bridge-group* **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 *bridge-group*, specify the bridge group number. The range is 1 to 255. You can create up to 32 bridge groups. Frames are bridged only among interfaces in the same group.
Step 3	**interface** *interface-id*	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: • 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** *vlan-id* 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** *bridge-group*	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.

	Command	Purpose
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. The **no bridge** *bridge-group* command automatically removes all SVIs and routes ports from that bridge group. 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 a port as a routed port, and to assign the port to the bridge group:

```
Switch(config)# bridge 10 protocol vlan-bridge
Switch(config)# interface gigabitethernet3/0/1
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 SVI for VLAN 2 and assigns it 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)# bridge-group 10
Switch(config-if)# exit
```

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 VLAN-Bridge Spanning-Tree Priority, page 54-6 (optional)
- Changing the Interface Priority, page 54-6 (optional)
- Assigning a Path Cost, page 54-7 (optional)
- Adjusting BPDU Intervals, page 54-7 (optional)
- Disabling the Spanning Tree on an Interface, page 54-9 (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 VLAN-Bridge Spanning-Tree Priority

You can globally configure the VLAN-bridge spanning-tree priority of a switch when it ties with another switch for the position as the root switch. You also can configure the likelihood that the switch will be selected as the root switch.

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** *bridge-group* **priority** *number*	Change the VLAN-bridge spanning-tree priority of the switch. • For *bridge-group*, specify the bridge group number. The range is 1 to 255. • For *number*, 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 a port, 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 a port. When two switches tie for position as the root switch, you configure a port 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** *interface-id*	Specify the interface to set the priority, and enter interface configuration mode.
Step 3	**bridge-group** *bridge-group* **priority** *number*	Change the priority of a port. • For *bridge-group*, specify the bridge group number. The range is 1 to 255. • For *number*, enter a number from 0 to 255 in increments of 4. The lower the number, the more likely that the port 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	**show running-config**	Verify your entry.
Step 6	**copy running-config startup-config**	(Optional) Save your entry in the configuration file.

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 to 20 on a port in bridge group 10:

```
Switch(config)# interface gigabitethernet2/0/1
Switch(config-if)# bridge-group 10 priority 20
```

Assigning a Path Cost

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

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

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Specify the port to set the path cost, and enter interface configuration mode.
Step 3	**bridge-group** *bridge-group* **path-cost** *cost*	Assign the path cost of a port. • For *bridge-group*, specify the bridge group number. The range is 1 to 255. • For *cost*, enter a number from 0 to 65535. The higher the value, the higher the cost. – For 10 Mb/s, the default path cost is 100. – For 100 Mb/s, the default path cost is 19. – For 1000 Mb/s, the default path cost is 4.
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 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 to 20 on a port in bridge group 10:

```
Switch(config)# interface gigabitethernet2/0/1
Switch(config-if)# bridge-group 10 path-cost 20
```

Adjusting BPDU Intervals

- Adjusting the Interval between Hello BPDUs, page 54-8 (optional)
- Changing the Forward-Delay Interval, page 54-8 (optional)
- Changing the Maximum-Idle Interval, page 54-9 (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** *bridge-group* **hello-time** *seconds*	Specify the interval between hello BPDUs. • For *bridge-group*, specify the bridge group number. The range is 1 to 255. • For *seconds*, enter a number from 1 to 10. The default is 2.
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 a port 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** *bridge-group* **forward-time** *seconds*	Specify the forward-delay interval. • For *bridge-group*, specify the bridge group number. The range is 1 to 255. • For *seconds*, enter a number from 4 to 200. The default is 20.
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** *bridge-group* **max-age** *seconds*	Specify the interval that the switch waits to hear BPDUs from the root switch. • For *bridge-group*, specify the bridge group number. The range is 1 to 255. • For *seconds*, enter a number from 6 to 200. The default is 30.
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 a port. This procedure is optional.

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Specify the port, and enter interface configuration mode.
Step 3	**bridge-group** *bridge-group* **spanning-disabled**	Disable spanning tree on the port. For *bridge-group*, specify the bridge group number. The range is 1 to 255.
Step 4	**end**	Return to privileged EXEC mode.

	Command	Purpose
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 port, use the **no bridge-group** *bridge-group* **spanning-disabled** interface configuration command.

This example shows how to disable spanning tree on a port in bridge group 10:

```
Switch(config)# interface gigabitethernet2/0/1
Switch(config-if)# bridge group 10 spanning-disabled
```

Monitoring and Maintaining Fallback Bridging

Table 54-2 Commands for Monitoring and Maintaining Fallback Bridging

Command	Purpose
clear bridge *bridge-group*	Removes any learned entries from the forwarding database.
show bridge [*bridge-group*] **group**	Displays details about the bridge group.
show bridge [*bridge-group*] [*interface-id* | *mac-address* | **verbose**	Displays MAC addresses learned in the bridge group.

To display the bridge-group MAC address table on a stack member, start a session from the stack master to the stack member by using the **session** *stack-member-number* global configuration command. Enter the **show bridge** [*bridge-group*] [*interface-id* | *mac-address* | **verbose**] privileged EXEC command at the stack member prompt.

For information about the fields in these displays, see the *Cisco IOS Bridging and IBM Networking Command Reference, Volume 1 of 2, Release 12.4.*