Spanning Tree Protocol

This document descibes Spanning Tree Protocol (STP) in a wireless environment.

Understanding Spanning Tree Protocol

This section describes how spanning-tree features work.

STP Overview

STP is a Layer 2 link management protocol that provides path redundancy while preventing loops in the network. For a Layer 2 network to function properly, only one active path can exist between any two stations. Spanning-tree operation is transparent to end stations, which cannot detect whether they are connected to a single LAN segment or to a LAN of multiple segments.

When you create fault-tolerant internetworks, you must have a loop-free path between all nodes in a network. The spanning-tree algorithm calculates the best loop-free path throughout a Layer 2 network. Infrastructure devices such as wireless bridges and switches send and receive spanning-tree frames, called *bridge protocol data units* (BPDUs), at regular intervals. The devices do not forward these frames but use them to construct a loop-free path.

Multiple active paths among end stations cause loops in the network. If a loop exists in the network, end stations might receive duplicate messages. Infrastructure devices might also learn end-station MAC addresses on multiple Layer 2 interfaces. These conditions result in an insecure network.

STP defines a tree with a root device and a loop-free path from the root to all infrastructure devices in the Layer 2 network.

Note In discussions about STP, the term *root* is used to describe two concepts: the bridge on the network that serves as a central point in the spanning tree is called the *root device*, and the port on each device that provides the most efficient path to the device is called the *root port*. These meanings are separate. A bridge whose role in radio network setting is root device does not necessarily become the root device in the spanning tree. In this chapter, the root device in the spanning tree is called the *spanning-tree root*.

Americas Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

STP forces redundant data paths into a standby (blocked) state. If a network segment in the spanning tree fails and a redundant path exists, the spanning-tree algorithm recalculates the spanning-tree topology and activates the standby path.

When two interfaces are part of a loop, the spanning-tree port priority and path cost settings determine which interface is put in the forwarding state and which is put in the blocking state. The port priority value represents the location of an interface in the network topology and how well it is located to pass traffic. The path cost value represents media speed.

The bridge supports both per-VLAN spanning tree (PVST) and a single 802.1q spanning tree without VLANs. The bridge cannot run 802.1s multiple spanning tree (MST) or 802.1d Common Spanning Tree, which map multiple VLANs into a one-instance spanning tree.

The bridge maintains a separate spanning-tree instance for each active VLAN configured on it. A bridge ID, consisting of the bridge priority and the MAC address, is associated with each instance. For each VLAN, the bridge with the lowest bridge ID becomes the spanning-tree root for that VLAN.

Bridge Interoperability

Cisco bridges are interoperable when STP is enabled and no VLANs are configured. This configuration is the only one possible, for the following reasons:

- When STP is disabled, the bridge acts as an access point and disallows association of non-root bridge.
- The bridge has a single instance of STP in non-VLAN configurations and multiple instances of STP in VLAN configurations.
- Incompatibilities between single and multiple instances of STP can cause inconsistent blocking of traffic when VLANs are configured. When the native VLAN is blocked, bridge flapping can occur.

Therefore, the best configuration for STP interoperability is to have the bridge STP feature enabled and VLANs not configured.

Note When Cisco bridges are configured as workgroup bridges, they can operate with STP disabled and allow for associations with access points. However, this configuration is not technically a bridge-to-bridge scenario.

Bridge Protocol Data Units

The stable, active spanning-tree topology of a network is determined by the following factors:

- The unique bridge ID (wireless bridge priority and MAC address) associated with each VLAN on each wireless bridge
- The spanning-tree path cost to the spanning-tree root
- The port identifier (port priority and MAC address) associated with each Layer 2 interface

When the bridges in a network are powered up, each bridge functions as the STP root. The bridges send configuration BPDUs through the Ethernet and radio ports. The BPDUs communicate and compute the spanning-tree topology. Each configuration BPDU contains this information:

- The unique bridge ID of the wireless bridge that the sending bridge identifies as the spanning-tree root
- The spanning-tree path cost to the root
- The bridge ID of the sending bridge
- Message age
- The identifier of the sending interface
- Values for the hello, forward delay, and max-age protocol timers

When a bridge receives a configuration BPDU that contains information *superior* (lower bridge ID, lower path cost, and so forth), it stores the information for that port. If this BPDU is received on the root port of the bridge, the bridge also forwards it with an updated message to all attached LANs for which it is the designated bridge.

If a bridge receives a configuration BPDU that contains *inferior* information to that currently stored for that port, it discards the BPDU. If the bridge is a designated bridge for the LAN from which the inferior BPDU was received, it sends that LAN a BPDU containing the up-to-date information stored for that port. In this way, inferior information is discarded, and superior information is propagated on the network.

A BPDU exchange results in these actions:

- One bridge is elected as the spanning-tree root.
- A root port is selected for each bridge (except the spanning-tree root). This port provides the best path (lowest cost) when the bridge forwards packets to the spanning-tree root.
- The shortest distance to the spanning-tree root is calculated for each bridge based on the path cost.
- A designated bridge for each LAN segment is selected. The designated bridge incurs the lowest path cost when forwarding packets from that LAN to the spanning-tree root. The port through which the designated bridge is attached to the LAN is called the *designated port*.
- Interfaces that are included in the spanning-tree instance are selected. Root ports and designated ports are put in the forwarding state.
- All interfaces that are not included in the spanning tree are blocked.

Election of the Spanning-Tree Root

All bridges in the Layer 2 network that are participating in STP gather information about other bridges in the network by exchanging BPDU data messages. This message exchange results in these actions:

- The election of a unique spanning-tree root for each spanning-tree instance
- The election of a designated bridge for every LAN segment
- The removal of loops in the network by blocking Layer 2 interfaces connected to redundant links

For each VLAN, the bridge with the highest bridge priority (the lowest numerical priority value) is elected as the spanning-tree root. If all bridges are configured with the default priority (32768), the bridge with the lowest MAC address in the VLAN becomes the spanning-tree root. The bridge priority value occupies the most significant bits of the bridge ID.

When you change the bridge priority value, you change the probability that the bridge will be elected as the root device. Configuring a higher value decreases the probability; a lower value increases the probability.

The spanning-tree root is the logical center of the spanning-tree topology. Paths that are not needed to reach the spanning-tree root from anywhere in the network are placed in the spanning-tree blocking mode.

BPDUs contain information about the sending bridge and its ports, including bridge and MAC addresses, bridge priority, port priority, and path cost. STP uses this information to elect the spanning-tree root and root port for the network and the root port and designated port for each LAN segment.

Spanning-Tree Timers

Table 1 describes the timers that affect the entire spanning-tree performance.

Table 1 Spanning-Tree Timers

Variable	Description
Hello timer	Determines how often the bridge broadcasts hello messages to other bridges.
Forward-delay timer	Determines how long each of the listening and learning states last before the interface begins forwarding.
Maximum-age timer	Determines the amount of time the bridge stores protocol information received on an interface.

Creating the Spanning-Tree Topology

In Figure 1, bridge 4 is elected as the spanning-tree root, under the assumption that the priority of all the bridges is set to the default (32768) and bridge 4 has the lowest MAC address. However, because of traffic patterns, number of forwarding interfaces, or link types, bridge 4 might not be the ideal spanning-tree root. By increasing the priority (lowering the numerical value) of the ideal bridge so that it becomes the spanning-tree root, you force a spanning-tree recalculation to form a new topology with the ideal bridge as the spanning-tree root.

Figure 1 Spanning-Tree Topology

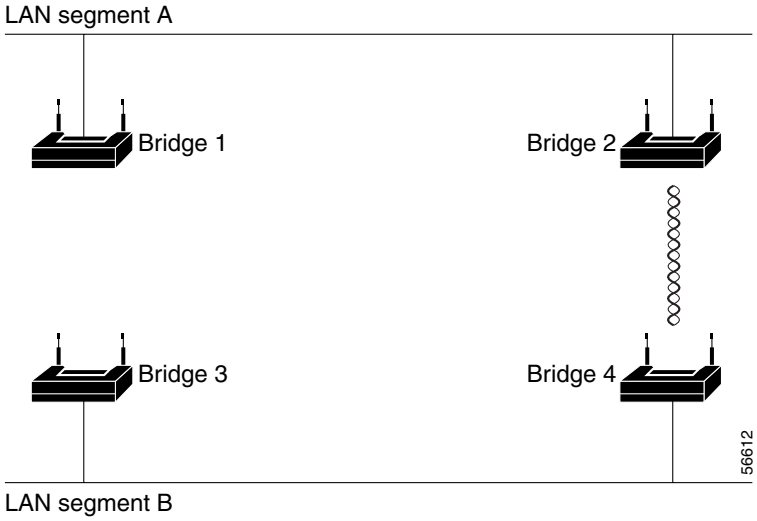

Spanning-Tree Interface States

Propagation delays can occur when protocol information passes through a wireless LAN. As a result, topology changes can take place at different times and at different places in the network. When an interface transitions directly from nonparticipation in the spanning-tree topology to the forwarding state, it can create temporary data loops. Interfaces must wait for new topology information to propagate through the LAN before starting to forward frames. They must allow the frame lifetime to expire for forwarded frames that have used the old topology.

Each interface on a bridge using spanning tree exists in one of these states:

- Blocking—The interface does not participate in frame forwarding.
- Listening—The first transitional state after the blocking state, in which the spanning tree determines that the interface should participate in frame forwarding.
- Learning—The interface prepares to participate in frame forwarding.
- Forwarding—The interface forwards frames.
- Disabled—The interface is not participating in spanning tree because there is a port shutdown, there is no link on the port, or no spanning-tree instance is running on the port.

An interface moves through these states:

- From initialization to blocking
- From blocking to listening or to disabled
- From listening to learning or to disabled
- From learning to forwarding or to disabled
- From forwarding to disabled

Figure 2 illustrates how an interface moves through the states.

Figure 2 Spanning-Tree Interface States

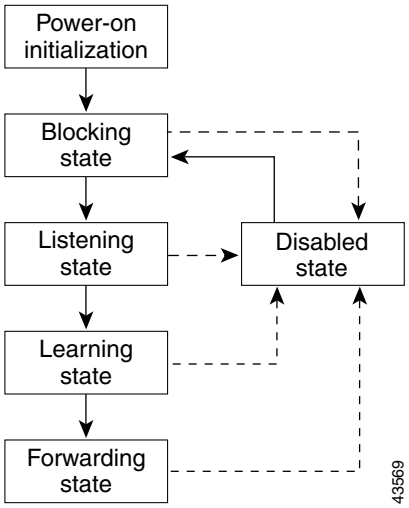

When you enable STP on the bridge, the Ethernet and radio interfaces go through the blocking state and the transitory states of listening and learning. STP stabilizes each interface at the forwarding or blocking state.

When the spanning-tree algorithm places a Layer 2 interface in the forwarding state, this process occurs:

1. The interface is in the listening state while STP waits for protocol information to transition the interface to the blocking state.
2. While STP waits the forward-delay timer to expire, it moves the interface to the learning state and resets the forward-delay timer.
3. In the learning state, the interface continues to block frame forwarding as the bridge learns end-station location information for the forwarding database.
4. When the forward-delay timer expires, STP moves the interface to the forwarding state, where both learning and frame forwarding are enabled.

Blocking State

An interface in the blocking state does not participate in frame forwarding. After initialization, a BPDU is sent to the bridge's Ethernet and radio ports. A bridge initially functions as the spanning-tree root until it exchanges BPDUs with other bridges. This exchange establishes which bridge in the network is the spanning-tree root. If there is only one bridge in the network, no exchange occurs, the forward-delay timer expires, and the interfaces move to the listening state. An interface always enters the blocking state when you enable STP.

An interface in the blocking state performs as follows:

- Discards frames received on the port
- Does not learn addresses
- Receives BPDUs

Note If a port is blocked, some broadcast or multicast packets can reach a forwarding port on the bridge and cause the bridging logic to switch the blocked port into listening state momentarily before the packets are dropped at the blocked port.

Listening State

The listening state is the first state an interface enters after the blocking state. The interface enters this state when STP determines that the interface should participate in frame forwarding.

An interface in the listening state performs as follows:

- Discards frames received on the port
- Does not learn addresses
- Receives BPDUs

Learning State

An interface in the learning state prepares to participate in frame forwarding. The interface enters the learning state from the listening state.

An interface in the learning state performs as follows:

- Discards frames received on the port
- Learns addresses
- Receives BPDUs

Forwarding State

An interface in the forwarding state forwards frames. The interface enters the forwarding state from the learning state.

An interface in the forwarding state performs as follows:

- Receives and forwards frames received on the port
- Learns addresses
- Receives BPDUs

Disabled State

An interface in the disabled state does not participate in frame forwarding or in the spanning tree. An interface in the disabled state is nonoperational.

A disabled interface performs as follows:

- Discards frames received on the port
- Does not learn addresses
- Does not receive BPDUs

Configuring STP Features

You complete three major steps to configure STP on the WMIC:

1. If necessary, assign interfaces and subinterfaces to bridge groups
2. Enable STP for each bridge group
3. Set the STP priority for each bridge group

These sections provide information on STP configuration:

Default STP Configuration

STP is disabled by default. Table 2 lists the default STP settings when you enable STP.

Table 2 Default STP Values When STP is Enabled

Setting	Default Value
bridge priority	32768
bridge max age	20
bridge hello time	2
bridge forward delay	15
Ethernet port path cost	19
Ethernet port priority	128
radio port path cost	33
radio port priority	128

The radio and Ethernet interfaces and the native VLAN on the bridge are assigned to bridge group 1 by default. When you enable STP and assign a priority on bridge group 1, STP is enabled on the radio and Ethernet interfaces and on the primary VLAN, and those interfaces adopt the priority assigned to bridge group 1. You can create bridge groups for subinterfaces and assign different STP settings to those bridge groups.

Configuring STP Settings

To configure STP on the WMIC, follow these steps, beginning in privileged EXEC mode:

	Command	Purpose
Step 1	**configure terminal**	Enters global configuration mode.
Step 2	**interface {dot11radio** *port* | **fastethernet** *port* }	Enters interface configuration mode for radio or Ethernet interfaces or subinterfaces.

	Command	Purpose
Step 3	**bridge-group** *number*	Assigns the interface to a bridge group. You can number your bridge groups from 1 to 255.
Step 4	**no bridge-group** *number* **spanning-disabled**	Counteracts the command that automatically disables STP for a bridge group. (STP is later enabled on the interface when you enter the bridge *n* **protocol ieee** command.)
Step 5	**exit**	Returns to global configuration mode.
Step 6	**bridge** *number* **protocol ieee**	Enables STP for the bridge group. You must enable STP on each bridge group that you create with **bridge-group** commands.
Step 7	**bridge** *number* **priority** *priority*	(Optional) Assigns a priority to a bridge group. The lower the priority, the more likely it is that the bridge becomes the spanning-tree root.
Step 8	**end**	Returns to privileged EXEC mode.
Step 9	**show spanning-tree bridge**	Verifies your entries.
Step 10	**copy running-config startup-config**	(Optional) Saves your entries in the configuration file.

STP Configuration Examples

These configuration examples show how to enable STP on root and non-root bridges with and without VLANs:

Root Device Without VLANs

This example shows the configuration of a root device with no VLANs configured and with STP enabled:

```
hostname master-bridge-south
ip subnet-zero
!
bridge irb
!
interface Dot11Radio0
no ip address
no ip route-cache
!
ssid tsunami
authentication open
guest-mode
!
speed basic-6.0 9.0 12.0 18.0 24.0 36.0 48.0 54.0
rts threshold 2312
station-role root
no cdp enable
infrastructure-client
bridge-group 1
```

```
!
interface FastEthernet0
no ip address
no ip route-cache
duplex auto
speed auto
bridge-group 1
!
interface BVI1
ip address 1.4.64.23 255.255.0.0
no ip route-cache
!
ip default-gateway 1.4.0.1
bridge 1 protocol ieee
bridge 1 route ip
bridge 1 priority 9000
!
line con 0
exec-timeout 0 0
line vty 0 4
login
line vty 5 15
login
!
end
```

Non-Root Bridge Without VLANs

This example shows the configuration of a non-root bridge with STP enabled and no VLANs configured:

```
hostname client-bridge-north
ip subnet-zero
!
bridge irb
!
interface Dot11Radio0
no ip address
no ip route-cache
!
ssid tsunami
authentication open
guest-mode
!
speed basic-6.0 9.0 12.0 18.0 24.0 36.0 48.0 54.0
rts threshold 2312
station-role non-root
no cdp enable
bridge-group 1
!
interface FastEthernet0
no ip address
no ip route-cache
duplex auto
speed auto
bridge-group 1 path-cost 40
!
interface BVI1
ip address 1.4.64.24 255.255.0.0
no ip route-cache
!
bridge 1 protocol ieee
```

```
bridge 1 route ip
bridge 1 priority 10000
!
line con 0
line vty 0 4
login
line vty 5 15
login
!
end
```

Root Device with VLANs

This example shows the configuration of a root device with VLANs configured with STP enabled:

```
hostname master-bridge-hq
!
ip subnet-zero
!
ip ssh time-out 120
ip ssh authentication-retries 3
!
bridge irb
!
interface Dot11Radio0
no ip address
no ip route-cache
!
ssid vlan1
vlan 1
infrastructure-ssid
authentication open
!
speed basic-6.0 9.0 12.0 18.0 24.0 36.0 48.0 54.0
rts threshold 2312
station-role root
no cdp enable
infrastructure-client
!
interface Dot11Radio0.1
encapsulation dot1Q 1 native
no ip route-cache
no cdp enable
bridge-group 1
!
interface Dot11Radio0.2
encapsulation dot1Q 2
no ip route-cache
no cdp enable
bridge-group 2
!
interface Dot11Radio0.3
encapsulation dot1Q 3
no ip route-cache
bridge-group 3
bridge-group 3 path-cost 500
!
interface FastEthernet0
no ip address
no ip route-cache
duplex auto
speed auto
!
```

```
interface FastEthernet0.1
encapsulation dot1Q 1 native
no ip route-cache
bridge-group 1
!
interface FastEthernet0.2
encapsulation dot1Q 2
no ip route-cache
bridge-group 2
!
interface FastEthernet0.3
encapsulation dot1Q 3
no ip route-cache
bridge-group 3
!
interface BVI1
ip address 1.4.64.23 255.255.0.0
no ip route-cache
!
ip default-gateway 1.4.0.1
bridge 1 protocol ieee
bridge 1 route ip
bridge 1 priority 9000
bridge 2 protocol ieee
bridge 2 priority 10000
bridge 3 protocol ieee
bridge 3 priority 3100
!
line con 0
exec-timeout 0 0
line vty 5 15
!
end
```

Non-Root Bridge with VLANs

This example shows the configuration of a non-root bridge with VLANs configured with STP enabled:

```
hostname client-bridge-remote
!
ip subnet-zero
!
ip ssh time-out 120
ip ssh authentication-retries 3
!
bridge irb
!
interface Dot11Radio0
no ip address
no ip route-cache
!
ssid vlan1
vlan 1
authentication open
infrastructure-ssid
!
speed basic-6.0 9.0 12.0 18.0 24.0 36.0 48.0 54.0
rts threshold 2312
station-role non-root
no cdp enable
!
interface Dot11Radio0.1
```

```
encapsulation dot1Q 1 native
no ip route-cache
no cdp enable
bridge-group 1
!
interface Dot11Radio0.2
encapsulation dot1Q 2
no ip route-cache
no cdp enable
bridge-group 2
!
interface Dot11Radio0.3
encapsulation dot1Q 3
no ip route-cache
no cdp enable
bridge-group 3
!
interface FastEthernet0
no ip address
no ip route-cache
duplex auto
speed auto
!
interface FastEthernet0.1
encapsulation dot1Q 1 native
no ip route-cache
bridge-group 1
!
interface FastEthernet0.2
encapsulation dot1Q 2
no ip route-cache
bridge-group 2
!
interface FastEthernet0.3
encapsulation dot1Q 3
no ip route-cache
bridge-group 3
bridge-group 3 path-cost 400
!
interface BVI1
ip address 1.4.64.24 255.255.0.0
no ip route-cache
!
bridge 1 protocol ieee
bridge 1 route ip
bridge 1 priority 10000
bridge 2 protocol ieee
bridge 2 priority 12000
bridge 3 protocol ieee
bridge 3 priority 2900
!
line con 0
line vty 5 15
!
end
```

Displaying Spanning-Tree Status

To display the spanning-tree status, use one or more of the commands in Table 3 in privileged EXEC mode:

Table 3 Commands for Displaying Spanning-Tree Status

Command	Purpose
show spanning-tree	Displays information on your network's spanning tree.
show spanning-tree blocked-ports	Displays a list of blocked ports on this device.
show spanning-tree bridge	Displays status and configuration of this bridge.
show spanning-tree active	Displays spanning-tree information on active interfaces only.
show spanning-tree root	Displays a detailed summary of information on the spanning-tree root.
show spanning-tree interface *interface-id*	Displays spanning-tree information for the specified interface.
show spanning-tree summary [**totals**]	Displays a summary of port states or displays the total lines of the STP state section.

For information about other keywords for the **show spanning-tree** privileged EXEC command, refer to the *Cisco IOS Command Reference for Cisco Access Points and Bridges*.