



Configuring Rapid PVST+ Using Cisco NX-OS

This chapter describes how to configure the Rapid per VLAN Spanning Tree (Rapid PVST+) protocol on Cisco NX-OS devices.

This chapter includes the following sections:

- [Finding Feature Information, on page 1](#)
- [Information About Rapid PVST+, on page 1](#)
- [Prerequisites for Configuring Rapid PVST+, on page 18](#)
- [Guidelines and Limitations for Configuring Rapid PVST+, on page 18](#)
- [Default Settings for Rapid PVST+, on page 19](#)
- [Configuring Rapid PVST+, on page 20](#)
- [Verifying the Rapid PVST+ Configurations, on page 36](#)
- [Displaying and Clearing Rapid PVST+ Statistics—CLI Version, on page 36](#)
- [Rapid PVST+ Example Configurations, on page 37](#)
- [Additional References for Rapid PVST+—CLI Version, on page 37](#)
- [Feature History for Configuring Rapid PVST+—CLI Version, on page 38](#)

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter or the Feature History table in this chapter.

Information About Rapid PVST+

**Note**

See the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide* for information on creating Layer 2 interfaces.

The Spanning Tree Protocol (STP) was implemented to provide a loop-free network at Layer 2 of the network. Rapid PVST+ is an updated implementation of STP that allows you to create one spanning tree topology for each VLAN. Rapid PVST+ is the default STP mode on the device.



Note Spanning tree is used to refer to IEEE 802.1w and IEEE 802.1s. If the IEEE 802.1D Spanning Tree Protocol is discussed in this publication, then 802.1D is stated specifically.



Note You can run either Rapid PVST+ or MST within each virtual device context (VDC). You cannot run both STP modes simultaneously in a VDC. Rapid PVST+ is the default STP mode.

The Rapid PVST+ protocol is the IEEE 802.1w standard, Rapid Spanning Tree Protocol (RSTP), implemented on a per VLAN basis. Rapid PVST+ interoperates with the IEEE 802.1Q VLAN standard, which mandates a single STP instance for all VLANs, rather than per VLAN.

Rapid PVST+ is enabled by default on the default VLAN (VLAN1) and on all newly created VLANs on the device. Rapid PVST+ interoperates with devices that run legacy IEEE 802.1D STP.

RSTP is an improvement on the original STP standard, 802.1D, which allows faster convergence.



Note The device supports full nondisruptive upgrades for Rapid PVST+. See the *Cisco Nexus 7000 Series NX-OS High Availability and Redundancy Guide* for complete information on nondisruptive upgrades.

The Cisco NX-OS release that is running on a managed device may not support all the features or settings described in this chapter. For the latest feature information and caveats, see the documentation and release notes for your platform and software release.



Note Beginning with Cisco NX-OS Release 5.x, when you are running virtual port channels (vPCs), you can configure STP for better performance. See the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide* for more information on this feature.

STP

STP is a Layer 2 link-management protocol that provides path redundancy while preventing loops in the network.

Overview of STP

In order for a Layer 2 Ethernet network to function properly, only one active path can exist between any two stations. STP operation is transparent to end stations, which cannot detect whether they are connected to a single LAN segment or a switched LAN of multiple segments.

When you create fault-tolerant internetworks, you must have a loop-free path between all nodes in a network. The STP algorithm calculates the best loop-free path throughout a switched Layer 2 network. Layer 2 LAN

ports send and receive STP frames, which are called Bridge Protocol Data Units (BPDUs), at regular intervals. Network devices do not forward these frames but use the frames to construct a loop-free path.

Multiple active paths between end stations cause loops in the network. If a loop exists in the network, end stations might receive duplicate messages and network devices might learn end station MAC addresses on multiple Layer 2 LAN ports.

STP defines a tree with a root bridge and a loop-free path from the root to all network devices in the Layer 2 network. STP forces redundant data paths into a blocked state. If a network segment in the spanning tree fails and a redundant path exists, the STP algorithm recalculates the spanning tree topology and activates the blocked path.

When two Layer 2 LAN ports on a network device are part of a loop, the STP port priority and port path-cost setting determine which port on the device is put in the forwarding state and which port is put in the blocking state. The STP port priority value is the efficiency with which that location allows the port to pass traffic. The STP port path-cost value is derived from the media speed.

How a Topology is Created

All devices in a LAN that participate in a spanning tree gather information about other switches in the network by exchanging BPDUs. This exchange of BPDUs results in the following actions:

- The system elects a unique root switch for the spanning tree network topology.
- The system elects a designated switch for each LAN segment.
- The system eliminates any loops in the switched network by placing redundant switch ports in a backup state; all paths that are not needed to reach the root device from anywhere in the switched network are placed in an STP-blocked state.

The topology on an active switched network is determined by the following:

- The unique device identifier Media Access Control (MAC) address of the device that is associated with each device
- The path cost to the root that is associated with each switch port
- The port identifier that is associated with each switch port

In a switched network, the root switch is the logical center of the spanning tree topology. STP uses BPDUs to elect the root switch and root port for the switched network.



Note The **mac-address bpdu source version 2** command enables STP to use the new Cisco MAC address (00:26:0b:xx:xx:xx) as the source address of BPDUs generated on vPC ports.

To apply this command, you must have identical configurations for both vPC peer switches or peers.

We strongly recommend that you disable EtherChannel Guard on the edge devices before using this command to minimize traffic disruption from STP inconsistencies. Reenable the EtherChannel Guard after updating on both peers.

Bridge ID

Each VLAN on each network device has a unique 64-bit bridge ID that consists of a bridge priority value, an extended system ID (IEEE 802.1t), and an STP MAC address allocation.

Bridge Priority Value

The bridge priority is a 4-bit value when the extended system ID is enabled.

You can only specify a device bridge ID (used by the spanning tree algorithm to determine the identity of the root bridge; the lowest number is preferred) as a multiple of 4096.



Note In this device, the extended system ID is always enabled; you cannot disable the extended system ID.

Extended System ID

The device always uses the 12-bit extended system ID.

Figure 1: Bridge ID with Extended System ID

This figure shows the 12-bit extended system ID field that is part of the bridge ID.



This table shows how the system ID extension combined with the bridge ID functions as the unique identifier for a VLAN.

Table 1: Bridge Priority Value and Extended System ID with the Extended System ID Enabled

Bridge Priority Value				Extended System ID (Set Equal to the VLAN ID)											
Bit 16	Bit 15	Bit 14	Bit 13	Bit 12	Bit 11	Bit 10	Bit 9	Bit 8	Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1
32768	16384	8192	4096	2048	1024	512	256	128	64	32	16	8	4	2	1

STP MAC Address Allocation



Note MAC address reduction is always enabled on the device.

Because MAC address reduction is always enabled on the device, you should also enable MAC address reduction on all other Layer 2 connected network devices to avoid undesirable root bridge election and spanning tree topology issues.

When MAC address reduction is enabled, the root bridge priority becomes a multiple of 4096 plus the VLAN ID. You can only specify a device bridge ID (used by the spanning tree algorithm to determine the identity of the root bridge; the lowest number is preferred) as a multiple of 4096. Only the following values are possible:

- 0
- 4096

- 8192
- 12288
- 16384
- 20480
- 24576
- 28672
- 32768
- 36864
- 40960
- 45056
- 49152
- 53248
- 57344
- 61440

STP uses the extended system ID plus a MAC address to make the bridge ID unique for each VLAN.

**Note**

If another bridge in the same spanning tree domain does not run the MAC address reduction feature, it could win the root bridge ownership because of the finer granularity in the selection of its bridge ID.

BPDUs

Network devices transmit BPDUs throughout the STP instance. Each network device sends configuration BPDUs to communicate and compute the spanning tree topology. Each configuration BPDU contains the following minimal information:

- The unique bridge ID of the network device that the transmitting network device believes to be the root bridge
- The STP path cost to the root
- The bridge ID of the transmitting bridge
- The message age
- The identifier of the transmitting port
- Values for the hello, forward delay, and max-age protocol timer
- Additional information for STP extension protocols

When a network device transmits a Rapid PVST+ BPDU frame, all network devices connected to the VLAN on which the frame is transmitted receive the BPDU. When a network device receives a BPDU, it does not

forward the frame but instead uses the information in the frame to calculate a BPDU. If the topology changes, the device initiates a BPDU exchange.

A BPDU exchange results in the following:

- One network device is elected as the root bridge.
- The shortest distance to the root bridge is calculated for each network device based on the path cost.
- A designated bridge for each LAN segment is selected. This network device is closest to the root bridge through which frames are forwarded to the root.
- A root port is elected. This port provides the best path from the bridge to the root bridge.
- Ports included in the spanning tree are selected.

Election of the Root Bridge

For each VLAN, the network device with the highest bridge ID (that is, the lowest numerical ID value) is elected as the root bridge. If all network devices are configured with the default priority (32768), the network device with the lowest MAC address in the VLAN becomes the root bridge. 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 device will be elected as the root bridge. Configuring a lower value increases the probability; a higher value decreases the probability.

The STP root bridge is the logical center of each spanning tree topology in a Layer 2 network. All paths that are not needed to reach the root bridge from anywhere in the Layer 2 network are placed in STP blocking mode.

BPDU contains information about the transmitting bridge and its ports, including bridge and MAC addresses, bridge priority, port priority, and path cost. STP uses this information to elect the root bridge for the STP instance, to elect the root port that leads to the root bridge, and to determine the designated port for each Layer 2 segment.

Creating the Spanning Tree Topology

By increasing the priority (lowering the numerical value) of the ideal network device so that it becomes the root bridge, you force an STP recalculation to form a new spanning tree topology with the ideal network device as the root.

Figure 2: Spanning Tree Topology

In this figure, switch A is elected as the root bridge because the bridge priority of all the network devices is set to the default (32768) and switch A has the lowest MAC address. However, due to traffic patterns, the number of forwarding ports or link types, switch A might not be the ideal root bridge.



When the spanning tree topology is calculated based on default parameters, the path between the source and destination end stations in a switched network might not be ideal. For instance, connecting higher-speed links to a port that has a higher number than the current root port can cause a root-port change. The goal is to make the fastest link the root port.

For example, assume that one port on switch B is a fiber-optic link, and another port on switch B (an unshielded twisted-pair [UTP] link) is the root port. Network traffic might be more efficient over the high-speed fiber-optic

link. By changing the STP port priority on the fiber-optic port to a higher priority (lower numerical value) than the root port, the fiber-optic port becomes the new root port.

Rapid PVST+

Rapid PVST+ is the default spanning tree mode for the software and is enabled by default on the default VLAN and all newly created VLANs.

A single instance, or topology, of RSTP runs on each configured VLAN, and each Rapid PVST+ instance on a VLAN has a single root device. You can enable and disable STP on a per-VLAN basis when you are running Rapid PVST+.

Overview of Rapid PVST+

Rapid PVST+ is the IEEE 802.1w (RSTP) standard implemented per VLAN. A single instance of STP runs on each configured VLAN (if you do not manually disable STP). Each Rapid PVST+ instance on a VLAN has a single root switch. You can enable and disable STP on a per-VLAN basis when you are running Rapid PVST+.



Note Rapid PVST+ is the default STP mode for the device.

Rapid PVST+ uses point-to-point wiring to provide rapid convergence of the spanning tree. The spanning tree reconfiguration can occur in less than 1 second with Rapid PVST+ (in contrast to 50 seconds with the default settings in the 802.1D STP). The device automatically checks the PVID.



Note Rapid PVST+ supports one STP instance for each VLAN.

Using Rapid PVST+, STP convergence occurs rapidly. By default, each designated port in the STP sends out a BPDU every 2 seconds. On a designated port in the topology, if hello messages are missed three consecutive times, or if the maximum age expires, the port immediately flushes all protocol information in the table. A port considers that it loses connectivity to its direct neighbor designated port if it misses three BPDUs or if the maximum age expires. This rapid aging of the protocol information allows quick failure detection.

Rapid PVST+ provides for rapid recovery of connectivity following the failure of a device, a device port, or a LAN. It provides rapid convergence for edge ports, new root ports, and ports connected through point-to-point links as follows:

- Edge ports—When you configure a port as an edge port on an RSTP device, the edge port immediately transitions to the forwarding state. (This immediate transition was previously a Cisco-proprietary feature named PortFast.) You should only configure ports that connect to a single end station as edge ports. Edge ports do not generate topology changes when the link changes.

Enter the **spanning-tree port type** interface configuration command to configure a port as an STP edge port.



Note We recommend that you configure all ports connected to a Layer 2 host as edge ports.

- Root port—If Rapid PVST+ selects a new root port, it blocks the old root port and immediately transitions the new root port to the forwarding state.
- Point-to-point links—If you connect a port to another port through a point-to-point link and the local port becomes a designated port, it negotiates a rapid transition with the other port by using the proposal-agreement handshake to ensure a loop-free topology.

Rapid PVST+ achieves rapid transition to the forwarding state only on edge ports and point-to-point links. Although the link type is configurable, the system automatically derives the link type information from the duplex setting of the port. Full-duplex ports are assumed to be point-to-point ports, while half-duplex ports are assumed to be shared ports.

Edge ports do not generate topology changes, but all other designated and root ports generate a topology change (TC) BPDU when they either fail to receive three consecutive BPDUs from the directly connected neighbor or the maximum age times out. At this point, the designated or root port sends a BPDU with the TC flag set. The BPDUs continue to set the TC flag as long as the TC While timer runs on that port. The value of the TC While timer is the value set for the hello time plus 1 second. The initial detector of the topology change immediately floods this information throughout the entire topology.

When Rapid PVST+ detects a topology change, the protocol does the following:

- Starts the TC While timer with a value equal to twice the hello time for all the nonedge root and designated ports, if necessary.
- Flushes the MAC addresses associated with all these ports.

The topology change notification floods quickly across the entire topology. The system flushes dynamic entries immediately on a per-port basis when it receives a topology change.



Note The TCA flag is used only when the device is interacting with devices that are running legacy 802.1D STP.

The proposal and agreement sequence then quickly propagates toward the edge of the network and quickly restores connectivity after a topology change.

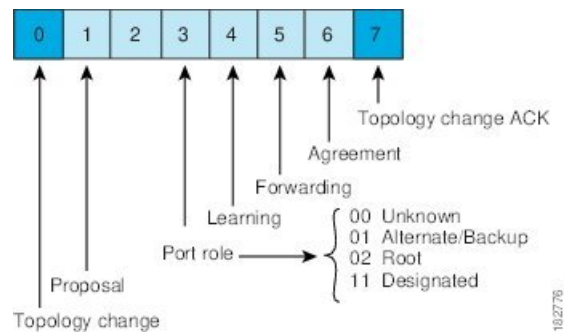
Rapid PVST+ BPDUs

Rapid PVST+ and 802.1w use all six bits of the flag byte to add the following:

- The role and state of the port that originates the BPDU
- The proposal and agreement handshake

Figure 3: Rapid PVST+ Flag Byte in BPDU

This figure shows the use of the BPDU flags in Rapid PVST+.



Another important change is that the Rapid PVST+ BPDU is type 2, version 2, which makes it possible for the device to detect connected legacy (802.1D) bridges. The BPDU for 802.1D is type 0, version 0.

Proposal and Agreement Handshake

Figure 4: Proposal and Agreement Handshaking for Rapid Convergence

In this figure, switch A is connected to switch B through a point-to-point link, and all of the ports are in the blocking state. Assume that the priority of switch A is a smaller numerical value than the priority of switch B. Switch A sends a proposal message (a configuration BPDU with the proposal flag set) to switch B, proposing itself as the designated switch.

After receiving the proposal message, switch B selects as its new root port the port from which the proposal message was received, forces all nonedge ports to the blocking state, and sends an agreement message (a BPDU with the agreement flag set) through its new root port.

After receiving the agreement message from switch B, switch A also immediately transitions its designated port to the forwarding state. No loops in the network can form because switch B blocked all of its nonedge ports and because there is a point-to-point link between switches A and B.

When switch C connects to switch B, a similar set of handshaking messages are exchanged. Switch C selects the port connected to switch B as its root port, and both ends of the link immediately transition to the forwarding state. With each iteration of this handshaking process, one more switch joins the active topology. As the network converges, this proposal-agreement handshaking progresses from the root toward the leaves of the spanning tree as shown in this figure.



The switch learns the link type from the port duplex mode; a full-duplex port is considered to have a point-to-point connection and a half-duplex port is considered to have a shared connection. You can override the default setting that is controlled by the duplex setting by entering the **spanning-tree link-type** interface configuration command.

This proposal/agreement handshake is initiated only when a nonedge port moves from the blocking to the forwarding state. The handshaking process then proliferates step-by-step throughout the topology.

Protocol Timers

This table describes the protocol timers that affect the Rapid PVST+ performance.

Table 2: Rapid PVST+ Protocol Timers

Variable	Description
Hello timer	Determines how often each device broadcasts BPDUs to other network devices. The default is 2 seconds, and the range is from 1 to 10.
Forward delay timer	Determines how long each of the listening and learning states last before the port begins forwarding. This timer is generally not used by the protocol, but it is used when interoperating with the 802.1D spanning tree. The default is 15 seconds, and the range is from 4 to 30 seconds.
Maximum age timer	Determines the amount of time that protocol information received on a port is stored by the network device. This timer is generally not used by the protocol, but it is used when interoperating with the 802.1D spanning tree. The default is 20 seconds; the range is from 6 to 40 seconds.

Port Roles


Rapid PVST+ provides rapid convergence of the spanning tree by assigning port roles and learning the active topology. Rapid PVST+ builds upon the 802.1D STP to select the device with the highest switch priority (lowest numerical priority value) as the root bridge. Rapid PVST+ assigns one of these port roles to individual ports:

- **Root port**—Provides the best path (lowest cost) when the device forwards packets to the root bridge.
- **Designated port**—Connects to the designated device that has the lowest path cost when forwarding packets from that LAN to the root bridge. The port through which the designated device is attached to the LAN is called the designated port.
- **Alternate port**—Offers an alternate path toward the root bridge to the path provided by the current root port. An alternate port provides a path to another device in the topology.
- **Backup port**—Acts as a backup for the path provided by a designated port toward the leaves of the spanning tree. A backup port can exist only when two ports are connected in a loopback by a point-to-point link or when a device has two or more connections to a shared LAN segment. A backup port provides another path in the topology to the device.
- **Disabled port**—Has no role within the operation of the spanning tree.

In a stable topology with consistent port roles throughout the network, Rapid PVST+ ensures that every root port and designated port immediately transition to the forwarding state while all alternate and backup ports are always in the blocking state. Designated ports start in the blocking state. The port state controls the operation of the forwarding and learning processes.

Figure 5: Sample Topology Demonstrating Port Roles

This figure shows port roles. A port with the root or a designated port role is included in the active topology. A port with the alternate or backup port role is excluded from the active topology.



182775

Rapid PVST+ Port State Overview

Propagation delays can occur when protocol information passes through a switched LAN. As a result, topology changes can take place at different times and at different places in a switched network. When a Layer 2 LAN port transitions directly from nonparticipation in the spanning tree topology to the forwarding state, it can create temporary data loops. Ports must wait for new topology information to propagate through the switched LAN before starting to forward frames.

Each Layer 2 LAN port on the device that uses Rapid PVST+ or MST exists in one of the following four states:

- Blocking—The Layer 2 LAN port does not participate in frame forwarding.
- Learning—The Layer 2 LAN port prepares to participate in frame forwarding.
- Forwarding—The Layer 2 LAN port forwards frames.
- Disabled—The Layer 2 LAN port does not participate in STP and is not forwarding frames.

When you enable Rapid PVST+, every port in the device, VLAN, and network goes through the blocking state and the transitory states of learning at power up. If properly configured, each Layer 2 LAN port stabilizes to the forwarding or blocking state.

When the STP algorithm places a Layer 2 LAN port in the forwarding state, the following process occurs:

1. The Layer 2 LAN port is put into the blocking state while it waits for protocol information that suggests it should go to the learning state.
2. The Layer 2 LAN port waits for the forward delay timer to expire, moves the Layer 2 LAN port to the learning state, and restarts the forward delay timer.
3. In the learning state, the Layer 2 LAN port continues to block frame forwarding as it learns the end station location information for the forwarding database.
4. The Layer 2 LAN port waits for the forward delay timer to expire and then moves the Layer 2 LAN port to the forwarding state, where both learning and frame forwarding are enabled.

Blocking State

A Layer 2 LAN port in the blocking state does not participate in frame forwarding.

A Layer 2 LAN port in the blocking state performs as follows:

- Discards frames received from the attached segment.
- Discards frames switched from another port for forwarding.
- Does not incorporate the end station location into its address database. (There is no learning on a blocking Layer 2 LAN port, so there is no address database update.)
- Receives BPDUs and directs them to the system module.
- Receives, processes, and transmits BPDUs received from the system module.
- Receives and responds to control plane messages.

Learning State

A Layer 2 LAN port in the learning state prepares to participate in frame forwarding by learning the MAC addresses for the frames. The Layer 2 LAN port enters the learning state from the blocking state.

A Layer 2 LAN port in the learning state performs as follows:

- Discards frames received from the attached segment.
- Discards frames switched from another port for forwarding.
- Incorporates the end station location into its address database.
- Receives BPDUs and directs them to the system module.
- Receives, processes, and transmits BPDUs received from the system module.
- Receives and responds to control plane messages.

Forwarding State

A Layer 2 LAN port in the forwarding state forwards frames. The Layer 2 LAN port enters the forwarding state from the learning state.

A Layer 2 LAN port in the forwarding state performs as follows:

- Forwards frames received from the attached segment.
- Forwards frames switched from another port for forwarding.
- Incorporates the end station location information into its address database.
- Receives BPDUs and directs them to the system module.
- Processes BPDUs received from the system module.
- Receives and responds to control plane messages.

Disabled State

A Layer 2 LAN port in the disabled state does not participate in frame forwarding or STP. A Layer 2 LAN port in the disabled state is virtually nonoperational.

A disabled Layer 2 LAN port performs as follows:

- Discards frames received from the attached segment.
- Discards frames switched from another port for forwarding.
- Does not incorporate the end station location into its address database. (There is no learning, so there is no address database update.)
- Does not receive BPDUs from neighbors.
- Does not receive BPDUs for transmission from the system module.

Summary of Port States

This table lists the possible operational and Rapid PVST+ states for ports and whether the port is included in the active topology.

Table 3: Port State Active Topology

Operational Status	Port State	Is Port Included in the Active Topology?
Enabled	Blocking	No
Enabled	Learning	Yes
Enabled	Forwarding	Yes
Disabled	Disabled	No

Synchronization of Port Roles

When the device receives a proposal message on one of its ports and that port is selected as the new root port, Rapid PVST+ forces all other ports to synchronize with the new root information.

The device is synchronized with superior root information received on the root port if all other ports are synchronized. An individual port on the device is synchronized if either of the following applies:

- That port is in the blocking state.
- It is an edge port (a port configured to be at the edge of the network).

If a designated port is in the forwarding state and is not configured as an edge port, it transitions to the blocking state when the Rapid PVST+ forces it to synchronize with new root information. In general, when the Rapid PVST+ forces a port to synchronize with root information and the port does not satisfy any of the above conditions, its port state is set to blocking.

After ensuring that all of the ports are synchronized, the device sends an agreement message to the designated device that corresponds to its root port. When the devices connected by a point-to-point link are in agreement about their port roles, Rapid PVST+ immediately transitions the port states to the forwarding state.

Figure 6: Sequence of Events During Rapid Convergence

This figure shows the sequence of events during synchronization.



Processing Superior BPDUs

A superior BPDU is a BPDU with root information (such as a lower switch ID or lower path cost) that is superior to what is currently stored for the port.

If a port receives a superior BPDU, Rapid PVST+ triggers a reconfiguration. If the port is proposed and is selected as the new root port, Rapid PVST+ forces all the designated, nonedge ports to synchronize.

If the received BPDU is a Rapid PVST+ BPDU with the proposal flag set, the device sends an agreement message after all of the other ports are synchronized. The new root port transitions to the forwarding state as soon as the previous port reaches the blocking state.

If the superior information received on the port causes the port to become a backup port or an alternate port, Rapid PVST+ sets the port to the blocking state and sends an agreement message. The designated port continues sending BPDUs with the proposal flag set until the forward-delay timer expires. At that time, the port transitions to the forwarding state.

Processing Inferior BPDUs

An inferior BPDU is a BPDU with root information (such as a higher switch ID or higher path cost) that is inferior to what is currently stored for the port.

If a designated port receives an inferior BPDU, it immediately replies with its own information.

Detecting Unidirectional Link Failure: Rapid PVST+

The software checks the consistency of the port role and state in the received BPDUs to detect unidirectional link failures that could cause bridging loops using the Unidirectional Link Detection (UDLD) feature. This feature is based on the dispute mechanism.

See the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide* for information on UDLD.

When a designated port detects a conflict, it keeps its role, but reverts to a discarding state because disrupting connectivity in case of inconsistency is preferable to opening a bridging loop.

Figure 7: Detecting Unidirectional Link Failure

This figure illustrates a unidirectional link failure that typically creates a bridging loop. Switch A is the root bridge, and its BPDUs are lost on the link leading to switch B. The 802.1w-standard BPDUs include the role and state of the sending port. With this information, switch A can detect that switch B does not react to the superior BPDUs that it sends and that switch B is the designated, not the root port. As a result, switch A blocks (or keeps blocking) its port, which prevents the bridging loop.



Port Cost



Note Rapid PVST+ uses the short (16-bit) path-cost method to calculate the cost by default. With the short path-cost method, you can assign any value in the range of 1 to 65535. However, you can configure the device to use the long (32-bit) path-cost method, which allows you to assign any value in the range of 1 to 200,000,000. You configure the path-cost calculation method globally.

This table shows how the STP port path-cost default value is determined from the media speed and path-cost calculation method of a LAN interface.

Table 4: Default Port Cost

Bandwidth	Short Path-Cost Method of Port Cost	Long Path-Cost Method of Port Cost
10 Mbps	100	2,000,000
100 Mbps	19	200,000
1 Gigabit Ethernet	4	20,000
10 Gigabit Ethernet	2	2,000

If a loop occurs, STP considers the port cost when selecting a LAN interface to put into the forwarding state.

You can assign the lower cost values to LAN interfaces that you want STP to select first and higher cost values to LAN interfaces that you want STP to select last. If all LAN interfaces have the same cost value, STP puts the LAN interface with the lowest LAN interface number in the forwarding state and blocks other LAN interfaces.

On access ports, you assign the port cost by the port. On trunk ports, you assign the port cost by the VLAN; you can configure the same port cost to all the VLANs on a trunk port.

Port Priority

If a redundant path occurs and multiple ports have the same path cost, Rapid PVST+ considers the port priority when selecting which LAN port to put into the forwarding state. You can assign lower priority values to LAN ports that you want Rapid PVST+ to select first and higher priority values to LAN ports that you want Rapid PVST+ to select last.

If all LAN ports have the same priority value, Rapid PVST+ puts the LAN port with the lowest LAN port number in the forwarding state and blocks other LAN ports. The possible priority range is from 0 through 224 (the default is 128), configurable in increments of 32. The device uses the port priority value when the LAN port is configured as an access port and uses the VLAN port priority values when the LAN port is configured as a trunk port.

Rapid PVST+ Scale Considerations

The following mentioned the scale considerations for Rapid PVST+.

Rapid PVST+ Logical Ports

Formula

Number of Logical Ports = Total Number of VLANs on all logical interfaces



Note Port-channel considered as one interface. Port-channel members are not counted individually.

In case of Rapid PVST+, VLAN Ports and Logical Ports are equal.

Rapid PVST+ Virtual Ports

Formula

Number of Virtual Ports = Total Number of VLANs on all physical interfaces



Note Port-channel members are counted individually.

For more information on verified scale limits, see [Cisco Nexus 7000 Series NX-OS Verified Scalability Guide](#).

Rapid PVST+ and IEEE 802.1Q Trunks

The 802.1Q trunks impose some limitations on the STP strategy for a network. In a network of Cisco network devices connected through 802.1Q trunks, the network devices maintain one instance of STP for each VLAN allowed on the trunks. However, non-Cisco 802.1Q network devices maintain only one instance of STP for all VLANs allowed on the trunks, which is the Common Spanning Tree (CST).

When you connect a Cisco network device to a non-Cisco device through an 802.1Q trunk, the Cisco network device combines the STP instance of the 802.1Q VLAN of the trunk with the STP instance of the non-Cisco 802.1Q network device. However, all per-VLAN STP information that is maintained by Cisco network devices is separated by a cloud of non-Cisco 802.1Q network devices. The non-Cisco 802.1Q cloud that separates the Cisco network devices is treated as a single trunk link between the network devices.

For more information on 802.1Q trunks, see the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide*.

Rapid PVST+ Interoperation with Legacy 802.1D STP

Rapid PVST+ can interoperate with devices that are running the legacy 802.1D protocol. The device knows that it is interoperating with equipment running 802.1D when it receives a BPDU version 0. The BPDUs for Rapid PVST+ are version 2. If the BPDU received is an 802.1w BPDU version 2 with the proposal flag set, the device sends an agreement message after all of the other ports are synchronized. If the BPDU is an 802.1D BPDU version 0, the device does not set the proposal flag and starts the forward-delay timer for the port. The new root port requires twice the forward-delay time to transition to the forwarding state.

The device interoperates with legacy 802.1D devices as follows:

- **Notification**—Unlike 802.1D BPDUs, 802.1w does not use TCN BPDUs. However, for interoperability with 802.1D devices, the device processes and generates TCN BPDUs.
- **Acknowledgment**—When an 802.1w device receives a TCN message on a designated port from an 802.1D device, it replies with an 802.1D configuration BPDU with the TCA bit set. However, if the TC-while timer (the same as the TC timer in 802.1D) is active on a root port connected to an 802.1D device and a configuration BPDU with the TCA set is received, the TC-while timer is reset.

This method of operation is required only for 802.1D devices. The 802.1w BPDUs do not have the TCA bit set.

- Protocol migration—For backward compatibility with 802.1D devices, 802.1w selectively sends 802.1D configuration BPDUs and TCN BPDUs on a per-port basis.

When a port is initialized, the migrate-delay timer is started (specifies the minimum time during which 802.1w BPDUs are sent), and 802.1w BPDUs are sent. While this timer is active, the device processes all BPDUs received on that port and ignores the protocol type.

If the device receives an 802.1D BPDU after the port migration-delay timer has expired, it assumes that it is connected to an 802.1D device and starts using only 802.1D BPDUs. However, if the 802.1w device is using 802.1D BPDUs on a port and receives an 802.1w BPDU after the timer has expired, it restarts the timer and starts using 802.1w BPDUs on that port.

**Note**

If you want all devices on the same LAN segment to reinitialize the protocol on each interface, you must reinitialize Rapid PVST+.

Rapid PVST+ Interoperation with 802.1s MST

Rapid PVST+ interoperates seamlessly with the IEEE 802.1s Multiple Spanning Tree (MST) standard. No user configuration is needed. To disable this seamless interoperation, you can use PVST Simulation.

High Availability for Rapid PVST+

The software supports high availability for Rapid PVST+. However, the statistics and timers are not restored when Rapid PVST+ restarts. The timers start again and the statistics begin from 0.

**Note**

See the *Cisco Nexus 7000 Series NX-OS High Availability and Redundancy Guide* for complete information on high-availability features.

Virtualization Support for Rapid PVST+

**Note**

See the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide* for complete information on virtual device contexts (VDCs) and assigning resources.

Figure 8: Separate STP in each VDC

This figure shows how the system provides support for VDCs. Using VDCs, you have a separate Layer 2 virtualization in each VDC, and each VDC runs a separate STP.

Each VDC will have its own Rapid PVST+. You cannot configure Rapid PVST+ across VDCs with Cisco NX-OS software. However, you can run Rapid PVST+ in one VDC and run MST in each VDC.

For example, VDC1 can run MST, VDC2 can run Rapid PVST+, and VDC3 can run MST.



Ensure that you are in the correct VDC before you begin configuring either Rapid PVST+ parameters.

Prerequisites for Configuring Rapid PVST+

Rapid PVST+ has the following prerequisites:

- You must be logged onto the device.
- If you are working in another VDC than the default VDC, that VDC must be created already.

Guidelines and Limitations for Configuring Rapid PVST+

Rapid PVST+ has the following configuration guidelines and limitations:

- There is a total of 4000 Rapid PVST+ for each VDC.
- The maximum number of VLANs and ports is 16,000.
- Only Rapid PVST+ or MST can be active at any time for each VDC.
- Port channeling—The port-channel bundle is considered as a single port. The port cost is the aggregation of all the configured port costs assigned to that channel.
- For private VLANs, on a normal VLAN trunk port, the primary and secondary private VLANs are two different logical ports and must have the exact STP topology. On access ports, STP sees only the primary VLAN.
- We recommend that you configure all ports connected to Layer 2 hosts as STP edge ports.
- Always leave STP enabled.
- Do not change timers because changing timers can adversely affect stability.
- Keep user traffic off the management VLAN; keep the management VLAN separate from the user data.
- Choose the distribution and core layers as the location of the primary and secondary root switches.
- When you connect two Cisco devices through 802.1Q trunks, the switches exchange spanning tree BPDUs on each VLAN allowed on the trunks. The BPDUs on the native VLAN of the trunk are sent untagged to the reserved 802.1D spanning tree multicast MAC address (01-80-C2-00-00-00). The BPDUs on all VLANs on the trunk are sent tagged to the reserved Cisco Shared Spanning Tree Protocol (SSTP) multicast MAC address (01-00-0c-cc-cc-cd).

Default Settings for Rapid PVST+

This table lists the default settings for Rapid PVST+ parameters.

Table 5: Default Rapid PVST+ Parameters

Parameters	Default
Spanning Tree	Enabled on all VLANs.
Spanning Tree mode	Rapid PVST+ Caution Changing the spanning tree mode disrupts the traffic because all spanning tree instances are stopped for the previous mode and started for the new mode.
VLAN	All ports assigned to VLAN1.
Extended system ID	Always enabled.
MAC address reduction	Always enabled.
Bridge ID priority	32769 (default bridge priority plus system ID extension of default VLAN1).
Port state	Blocking (changes immediately after convergence).
Port role	Designated (changes after convergence).
Port/VLAN priority	128.
Path-cost calculation method	Short.

Parameters	Default
Port/VLAN cost	<p>Auto</p> <p>The default port cost is determined by the media speed and path-cost method calculation, as follows:</p> <ul style="list-style-type: none"> • 10 Mbps: <ul style="list-style-type: none"> • short: 100 • long: 2,000,000 • 100 Mbps: <ul style="list-style-type: none"> • short: 19 • long: 200,000 • 1 Gigabit Ethernet: <ul style="list-style-type: none"> • short: 4 • long: 20,000 • 10 Gigabit Ethernet: <ul style="list-style-type: none"> • short: 2 • long: 2,000
Hello time	2 seconds.
Forward delay time	15 seconds.
Maximum aging time	20 seconds.
Link type	<p>Auto</p> <p>The default link type is determined by the duplex, as follows:</p> <ul style="list-style-type: none"> • Full duplex: point-to-point link • Half duplex: shared link

Configuring Rapid PVST+

Rapid PVST+, which has the 802.1w standard applied to the PVST+ protocol, is the default STP setting in the device.

You enable Rapid PVST+ on a per-VLAN basis. The device maintains a separate instance of STP for each VLAN (except on those VLANs on which you disable STP). By default, Rapid PVST+ is enabled on the default VLAN and on each VLAN that you create.



Note If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

Guidelines for Configuring Rapid PVST+

Ensure that you are in the correct VDC (see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide*).

Enabling Rapid PVST+—CLI Version

If you disable Rapid PVST+ on any VLANs, you must reenable Rapid PVRST+ on the specified VLANs. If you have enabled MST on the device and now want to use Rapid PVST+, you must enable Rapid PVST+ on the device.

Rapid PVST+ is the default STP mode. You cannot simultaneously run MST and Rapid PVST+ in the same VDC.

However, one VDC can run Rapid PVST+ and a different VDC can run MST.



Note When you change the spanning tree mode, traffic is disrupted because all spanning tree instances are stopped for the previous mode and started for the new mode.

Before you begin

Ensure that you are in the correct VDC (or enter the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	config t Example: switch# config t switch(config)#	Enters global configuration mode.
Step 2	spanning-tree mode rapid-pvst Example: switch(config)# spanning-tree mode rapid-pvst	Enables Rapid PVST+ on the device. Rapid PVST+ is the default spanning tree mode. Note Changing the spanning tree mode disrupts traffic because all spanning tree instances are stopped for the previous mode and started for the new mode.
Step 3	exit Example:	Exits global configuration mode.

	Command or Action	Purpose
	switch(config)# exit switch#	
Step 4	(Optional) show running-config spanning-tree all Example: switch# show running-config spanning-tree all	Displays information about the currently running STP configuration.
Step 5	(Optional) copy running-config startup-config Example: switch# copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to enable Rapid PVST+ on the device:

```
switch# config t
switch(config)# spanning-tree mode rapid-pvst
switch(config)# exit
switch#
```



Note Because Rapid PVST+ is enabled by default, entering the **show running** command to view the resulting configuration does not display the command that you entered to enable Rapid PVST+.

Disabling or Enabling Rapid PVST+ Per VLAN—CLI Version

You can enable or disable Rapid PVST+ on each VLAN.



Note Rapid PVST+ is enabled by default on the default VLAN and on all VLANs that you create.

Before you begin

Ensure that you are in the correct VDC (or enter the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	config t Example:	Enters global configuration mode.

	Command or Action	Purpose
	switch# config t switch(config)#	
Step 2	Enter one of the following commands:	
	Option	Description
	spanning-tree vlan vlan-range	Enables Rapid PVST+ (default STP) on a per VLAN basis. The <i>vlan-range</i> value can be 2 through 4094 except for reserved VLAN values.
	no spanning-tree vlan vlan-range	Disables Rapid PVST+ on the specified VLAN. See the Caution for information regarding this command.
	Example: switch(config)# spanning-tree vlan 5	
Step 3	exit Example: switch(config)# exit switch#	Exits global configuration mode.
Step 4	(Optional) show spanning-tree Example: switch# show spanning-tree	Displays the STP configuration.
Step 5	(Optional) copy running-config startup-config Example: switch# copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to enable STP on VLAN 5:

```
switch# config t
switch(config)# spanning-tree vlan 5
switch(config)# exit
switch#
```



Note Do not disable spanning tree on a VLAN unless all switches and bridges in the VLAN have spanning tree disabled. You cannot disable spanning tree on some switches and bridges in a VLAN and leave it enabled on other switches and bridges in the VLAN. This action can have unexpected results because switches and bridges with spanning tree enabled will have incomplete information regarding the physical topology of the network.



Caution We do not recommend disabling spanning tree even in a topology that is free of physical loops. Spanning tree serves as a safeguard against misconfigurations and cabling errors. Do not disable spanning tree in a VLAN without ensuring that no physical loops are present in the VLAN.



Note Because STP is enabled by default, entering the **show running** command to view the resulting configuration does not display the command that you entered to enable STP.

Configuring the Root Bridge ID

The device maintains a separate instance of STP for each active VLAN in Rapid PVST+. For each VLAN, the network device with the lowest bridge ID becomes the root bridge for that VLAN.

To configure a VLAN instance to become the root bridge, modify the bridge priority from the default value (32768) to a significantly lower value.

When you enter the **spanning-tree vlan *vlan_ID* primary root** command, the device checks the bridge priority of the current root bridges for each VLAN. The device sets the bridge priority for the specified VLANs to 24576 if this value will cause the device to become the root for the specified VLANs. If any root bridge for the specified VLAN has a bridge priority lower than 24576, the device sets the bridge priority for the specified VLANs to 4096 less than the lowest bridge priority.



Note The **spanning-tree vlan *vlan_ID* primary root** command fails if the value required to be the root bridge is less than 4096. If the software cannot lower the bridge priority any lower, the device returns the following message:

```
Error: Failed to set root bridge for VLAN 1
It may be possible to make the bridge root by setting the priority
for some (or all) of these instances to zero.
```



Caution The root bridge for each instance of STP should be a backbone or distribution device. Do not configure an access device as the STP primary root.

Enter the **diameter** keyword to specify the Layer 2 network diameter (that is, the maximum number of bridge hops between any two end stations in the Layer 2 network). When you specify the network diameter, the software automatically selects an optimal hello time, forward delay time, and maximum age time for a network

of that diameter, which can significantly reduce the STP convergence time. You can enter the **hello-time** keyword to override the automatically calculated hello time.



Note With the device configured as the root bridge, do not manually configure the hello time, forward-delay time, and maximum-age time using the **spanning-tree mst hello-time**, **spanning-tree mst forward-time**, and **spanning-tree mst max-age** global configuration commands.

Before you begin

Ensure that you are in the correct VDC (or enter the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	config t Example: switch# config t switch(config)#	Enters global configuration mode.
Step 2	spanning-tree vlan <i>vlan-range</i> root primary [<i>diameter dia</i> [<i>hello-time hello-time</i>]] Example: switch(config)# spanning-tree vlan 5 root primary diameter 4	Configures a device as the primary root bridge. The <i>vlan-range</i> value can be 2 through 4094 (except for reserved VLAN values.) The <i>dia</i> default is 7. The <i>hello-time</i> can be from 1 to 10 seconds, and the default value is 2 seconds.
Step 3	exit Example: switch(config)# exit switch#	Exits global configuration mode.
Step 4	(Optional) show spanning-tree Example: switch# show spanning-tree	Displays the STP configuration.
Step 5	(Optional) copy running-config startup-config Example: switch# copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to configure the device as the root bridge for VLAN 5 with a network diameter of 4:

```
switch# config t
```

```
switch(config)# spanning-tree vlan 5 root primary
diameter 4
switch(config)# exit
switch#
```

Configuring a Secondary Root Bridge—CLI Version

When you configure a device as the secondary root, the STP bridge priority is modified from the default value (32768) so that the device is likely to become the root bridge for the specified VLANs if the primary root bridge fails (assuming the other network devices in the network use the default bridge priority of 32768). STP sets the bridge priority to 28672.

Enter the **diameter** keyword to specify the Layer 2 network diameter (that is, the maximum number of bridge hops between any two end stations in the Layer 2 network). When you specify the network diameter, the software automatically selects an optimal hello time, forward delay time, and maximum age time for a network of that diameter, which can significantly reduce the STP convergence time. You can enter the **hello-time** keyword to override the automatically calculated hello time.

You can configure more than one device in this manner to have multiple backup root bridges. Enter the same network diameter and hello time values that you used when configuring the primary root bridge.



Note With the device configured as the root bridge, do not manually configure the hello time, forward-delay time, and maximum-age time using the **spanning-tree mst hello-time**, **spanning-tree mst forward-time**, and **spanning-tree mst max-age** global configuration commands.

Before you begin

Ensure that you are in the correct VDC (or enter the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	config t Example: switch# config t switch(config)#	Enters global configuration mode.
Step 2	spanning-tree vlan <i>vlan-range</i> root secondary [<i>diameter dia</i> [<i>hello-time hello-time</i>]] Example: switch(config)# spanning-tree vlan 5 root secondary diameter 4	Configures a device as the secondary root bridge. The <i>vlan-range</i> value can be 2 through 4094 (except for reserved VLAN values). The <i>dia</i> default is 7. The <i>hello-time</i> can be from 1 to 10 seconds, and the default value is 2 seconds.
Step 3	exit Example: switch(config)# exit switch#	Exits global configuration mode.

	Command or Action	Purpose
Step 4	(Optional) show spanning-tree vlan <i>vlan_id</i> Example: switch# show spanning-tree vlan 5	Displays the STP configuration for the specified VLANs.
Step 5	(Optional) copy running-config startup-config Example: switch# copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to configure the device as the secondary root bridge for VLAN 5 with a network diameter of 4:

```
switch# config t
switch(config)# spanning-tree vlan 5 root secondary diameter 4
switch(config)# exit
switch#
```

Configuring the Rapid PVST+ Bridge Priority of a VLAN

You can configure the Rapid PVST+ bridge priority of a VLAN. This is another method of configuring root bridges.



Note Be careful when using this configuration. We recommend that you configure the primary root and secondary root to modify the bridge priority.

Before you begin

Ensure that you are in the correct VDC (or enter the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	config t Example: switch# config t switch(config)#	Enters global configuration mode.
Step 2	spanning-tree vlan <i>vlan-range</i> priority <i>value</i> Example: switch(config)# spanning-tree vlan 5 priority 8192	Configures the bridge priority of a VLAN. Valid values are 0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, and 61440. All other values are rejected. The default value is 32768.

	Command or Action	Purpose
Step 3	exit Example: <pre>switch(config)# exit switch#</pre>	Exits global configuration mode.
Step 4	(Optional) show spanning-tree vlan <i>vlan_id</i> Example: <pre>switch# show spanning-tree vlan 5</pre>	Displays the STP configuration for the specified VLANs.
Step 5	(Optional) copy running-config startup-config Example: <pre>switch# copy running-config startup-config</pre>	Copies the running configuration to the startup configuration.

Example

This example shows how to configure the priority of VLAN 5 on Gigabit Ethernet port 1/4 to 8192:

```
switch# config t
switch(config)# spanning-tree vlan 5 priority 8192
switch(config)# exit
switch#
```

Configuring the Rapid PVST+ Port Priority—CLI Version

You can assign lower priority values to LAN ports that you want Rapid PVST+ to select first and higher priority values to LAN ports that you want Rapid PVST+ to select last. If all LAN ports have the same priority value, Rapid PVST+ puts the LAN port with the lowest LAN port number in the forwarding state and blocks other LAN ports.

The device uses the port priority value when the LAN port is configured as an access port and uses the VLAN port priority values when the LAN port is configured as a trunk port.

Before you begin

Ensure that you are in the correct VDC (or enter the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	config t Example: <pre>switch# config t switch(config)#</pre>	Enters global configuration mode.

	Command or Action	Purpose
Step 2	interface <i>type slot/port</i> Example: <pre>switch(config)# interface ethernet 1/4 switch(config-if)#</pre>	Specifies the interface to configure and enters interface configuration mode.
Step 3	spanning-tree [vlan <i>vlan-list</i>] port-priority <i>priority</i> Example: <pre>switch(config-if)# spanning-tree port-priority 160</pre>	Configures the port priority for the LAN interface. The <i>priority</i> value can be from 0 to 224. A lower value indicates a higher priority. The priority values are 0, 32, 64, 96, 128, 160, 192, and 224. All other values are rejected. The default value is 128.
Step 4	exit Example: <pre>switch(config-if)# exit switch(config)#</pre>	Exits interface configuration mode.
Step 5	(Optional) show spanning-tree interface {ethernet <i>slot/port</i> port channel <i>channel-number</i>} Example: <pre>switch# show spanning-tree interface ethernet 2/10</pre>	Displays the STP configuration for the specified interface.
Step 6	(Optional) copy running-config startup-config Example: <pre>switch(config)# copy running-config startup-config</pre>	Copies the running configuration to the startup configuration.

Example

This example shows how to configure the port priority of Ethernet access port 1/4 to 160:

```
switch# config t
switch (config)# interface ethernet 1/4
switch(config-if)# spanning-tree port-priority 160
switch(config-if)# exit
switch(config)#
```

Configuring the Rapid PVST+ Path-Cost Method and Port Cost—CLI Version

On access ports, you can assign the port cost for each port. On trunk ports, you can assign the port cost for each VLAN; you can configure all the VLANs on a trunk with the same port cost.



Note

In Rapid PVST+ mode, you can use either the short or long path-cost method, and you can configure the method in either the interface or configuration submenu. The default path-cost method is short.

Before you begin

Ensure that you are in the correct VDC (or enter the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	config t Example: <pre>switch# config t switch(config)#</pre>	Enters global configuration mode.
Step 2	spanning-tree pathcost method {long short} Example: <pre>switch(config)# spanning-tree pathcost method long</pre>	Selects the method used for Rapid PVST+ path-cost calculations. The default method is the short method.
Step 3	interface type slot/port Example: <pre>switch(config)# interface ethernet 1/4 switch(config-if)</pre>	Specifies the interface to configure and enters the interface configuration mode.
Step 4	spanning-tree [vlan vlan-id] cost [value auto] Example: <pre>switch(config-if)# spanning-tree cost 1000</pre>	<p>Configures the port cost for the LAN interface. The cost value, depending on the path-cost calculation method, can be as follows:</p> <ul style="list-style-type: none"> • short—1 to 65535 • long—1 to 200000000 <p>Note You configure this parameter per port on access ports and per VLAN on trunk ports.</p> <p>The default is auto, which sets the port cost on both the path-cost calculation method and the media speed.</p>
Step 5	exit Example: <pre>switch(config-if)# exit switch(config)#</pre>	Exits interface configuration mode.
Step 6	(Optional) show spanning-tree pathcost method Example: <pre>switch# show spanning-tree pathcost method</pre>	Displays the STP path-cost method.

	Command or Action	Purpose
Step 7	(Optional) copy running-config startup-config Example: <pre>switch(config)# copy running-config startup-config</pre>	Copies the running configuration to the startup configuration.

Example

This example shows how to configure the port cost of Ethernet access port 1/4 to 1000:

```
switch# config t
switch (config)# spanning-tree pathcost method long
switch (config)# interface ethernet 1/4
switch(config-if)# spanning-tree cost 1000
switch(config-if)# exit
switch(config)#
```

Configuring the Rapid PVST+ Hello Time for a VLAN—CLI Version

You can configure the Rapid-PVST+ hello time for a VLAN.



Note Be careful when using this configuration because you might disrupt the Spanning Tree. For most situations, we recommend that you configure the primary root and secondary root to modify the hello time.

Before you begin

Ensure that you are in the correct VDC (or enter the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	config t Example: <pre>switch# config t switch(config)#</pre>	Enters global configuration mode.
Step 2	spanning-tree vlan <i>vlan-range</i> hello-time <i>value</i> Example: <pre>switch(config)# spanning-tree vlan 5 hello-time 7</pre>	Configures the hello time of a VLAN. The hello time value can be from 1 to 10 seconds, and the default is 2 seconds.
Step 3	exit Example:	Exits global configuration mode.

	Command or Action	Purpose
	switch(config)# exit switch#	
Step 4	(Optional) show spanning-tree vlan <i>vlan_id</i> Example: switch# show spanning-tree vlan 5	Displays the STP configuration per VLAN.
Step 5	(Optional) copy running-config startup-config Example: switch# copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to configure the hello time for VLAN 5 to 7 seconds:

```
switch# config t
switch(config)# spanning-tree vlan 5 hello-time 7
switch(config)# exit
switch#
```

Configuring the Rapid PVST+ Forward Delay Time for a VLAN—CLI Version

You can configure the forward delay time per VLAN when using Rapid PVST+.

Before you begin

Ensure that you are in the correct VDC (or enter the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	config t Example: switch# config t switch(config)#	Enters global configuration mode.
Step 2	spanning-tree vlan <i>vlan-range</i> forward-time <i>value</i> Example: switch(config)# spanning-tree vlan 5 forward-time 21	Configures the forward delay time of a VLAN. The forward delay time value can be from 4 to 30 seconds, and the default is 15 seconds.
Step 3	exit Example:	Exits global configuration mode.

	Command or Action	Purpose
	switch(config)# exit switch#	
Step 4	(Optional) show spanning-tree vlan <i>vlan_id</i> Example: switch# show spanning-tree vlan 5	Displays the STP configuration per VLAN.
Step 5	(Optional) copy running-config startup-config Example: switch# copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to configure the forward delay time for VLAN 5 to 21 seconds:

```
switch# config t
switch(config)# spanning-tree vlan 5 forward-time 21
switch(config)# exit
switch#
```

Configuring the Rapid PVST+ Maximum Age Time for a VLAN—CLI Version

You can configure the maximum age time per VLAN when using Rapid PVST+.

Before you begin

Ensure that you are in the correct VDC (or enter the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	config t Example: switch# config t switch(config)#	Enters global configuration mode.
Step 2	spanning-tree vlan <i>vlan-range max-age value</i> Example: switch(config)# spanning-tree vlan 5 max-age 36	Configures the maximum aging time of a VLAN. The maximum aging time value can be from 6 to 40 seconds, and the default is 20 seconds.
Step 3	exit Example: switch(config)# exit switch#	Exits global configuration mode.

	Command or Action	Purpose
Step 4	(Optional) show spanning-tree vlan <i>vlan_id</i> Example: switch# show spanning-tree vlan 5	Displays the STP configuration per VLAN.
Step 5	(Optional) copy running-config startup-config Example: switch# copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to configure the maximum aging time for VLAN 5 to 36 seconds:

```
switch# config t
switch(config)# spanning-tree vlan 5 max-age 36
switch(config)# exit
switch#
```

Specifying the Link Type for Rapid PVST+—CLI Version

Rapid connectivity (802.1w standard) is established only on point-to-point links. By default, the link type is controlled from the duplex mode of the interface. A full-duplex port is considered to have a point-to-point connection; a half-duplex port is considered to have a shared connection.

If you have a half-duplex link physically connected point to point to a single port on a remote device, you can override the default setting on the link type and enable rapid transitions.

If you set the link to shared, STP falls back to 802.1D.

Before you begin

Ensure that you are in the correct VDC (or enter the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	config t Example: switch# config t switch(config)#	Enters global configuration mode.
Step 2	interface <i>type slot/port</i> Example: switch(config)# interface ethernet 1/4 switch(config-if)#	Specifies the interface to configure and enters interface configuration mode.

	Command or Action	Purpose
Step 3	spanning-tree link-type { <i>auto</i> <i>point-to-point</i> <i>shared</i> } Example: <pre>switch(config-if)# spanning-tree link-type point-to-point</pre>	Configures the link type to be either a point-to-point link or shared link. The system reads the default value from the device connection, as follows: half duplex links are shared and full-duplex links are point to point. If the link type is shared, the STP falls back to 802.1D. The default is auto, which sets the link type based on the duplex setting of the interface.
Step 4	exit Example: <pre>switch(config-if)# exit switch(config)#</pre>	Exits interface configuration mode.
Step 5	(Optional) show spanning-tree Example: <pre>switch# show spanning-tree</pre>	Displays the STP configuration.
Step 6	(Optional) copy running-config startup-config Example: <pre>switch(config)# copy running-config startup-config</pre>	Copies the running configuration to the startup configuration.

Example

This example shows how to configure the link type as a point-to-point link:

```
switch# config t
switch (config)# interface ethernet 1/4
switch(config-if)# spanning-tree link-type point-to-point
switch(config-if)# exit
switch(config)#
```

Reinitializing the Protocol for Rapid PVST+

A bridge that runs Rapid PVST+ can send 802.1D BPDUs on one of its ports when it is connected to a legacy bridge. However, the STP protocol migration cannot determine whether the legacy device has been removed from the link unless the legacy device is the designated switch. You can reinitialize the protocol negotiation (force the renegotiation with neighboring devices) on the entire device or on specified interfaces.

Before you begin

Ensure that you are in the correct VDC (or enter the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	clear spanning-tree detected-protocol [interface {ethernet slot/port port channel channel-number}] Example: switch# clear spanning-tree detected-protocol	Reinitializes Rapid PVST+ on all interfaces on the device or specified interfaces.

Example

This example shows how to reinitialize Rapid PVST+ on the Ethernet interface on slot 2, port 8:

```
switch# clear spanning-tree detected-protocol interface ethernet 2/8
switch#
```

Verifying the Rapid PVST+ Configurations

To display Rapid PVST+ configuration information, perform one of the following tasks:

Command	Purpose
show running-config spanning-tree [all]	Displays STP information.
show spanning-tree summary	Displays summary STP information.
show spanning-tree detail	Displays detailed STP information.
show spanning-tree {vlan vlan-id interface {[ethernet slot/port] [port-channel channel-number]} } [detail]	Displays STP information per VLAN and interface.
show spanning-tree vlan vlan-id bridge	Displays information on the STP bridge.

For information about the fields in the output from these commands, see the *Cisco Nexus 7000 Series NX-OS Layer 2 Switching Command Reference*.

Displaying and Clearing Rapid PVST+ Statistics—CLI Version

To display Rapid PVRST+ configuration information, perform one of the following tasks:

Command	Purpose
clear spanning-tree counters [interface type slot/port vlan vlan-id]	Clears the counters for STP.
show spanning-tree {vlan vlan-id interface {[ethernet slot/port] [port-channel channel-number]} } detail	Displays information about STP by interface or VLAN including BPDUs sent and received.

Rapid PVST+ Example Configurations

The following example shows how to configure Rapid PVST+:

```
switch# configure terminal
switch(config)# spanning-tree port type edge bpduguard default
switch(config)# spanning-tree port type edge bpdufilter default
switch(config)# spanning-tree port type network default
switch(config)# spanning-tree vlan 1-10 priority 24576
switch(config)# spanning-tree vlan 1-10 hello-time 1
switch(config)# spanning-tree vlan 1-10 forward-time 9
switch(config)# spanning-tree vlan 1-10 max-age 13

switch(config)# interface Ethernet 3/1 switchport
switch(config-if)# spanning-tree port type edge
switch(config-if)# exit

switch(config)# spanning-tree port type edge
switch(config-if)# switchport
switch(config-if)# switchport mode trunk
switch(config-if)# spanning-tree guard root
switch(config-if)# exit
switch(config)#
```

Additional References for Rapid PVST+—CLI Version

Related Documents

Related Topic	Document Title
Command reference	<i>Cisco Nexus 7000 Series NX-OS Layer 2 Switching Command Reference</i>
Layer 2 interfaces	<i>Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide</i>
NX-OS fundamentals	<i>Cisco Nexus 7000 Series NX-OS Fundamentals Configuration Guide</i>
High availability	<i>Cisco Nexus 7000 Series NX-OS High Availability and Redundancy Guide</i>
System management	<i>Cisco Nexus 7000 Series NX-OS System Management Configuration Guide</i>
VDCs	<i>Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide</i>
Licensing	<i>Cisco NX-OS Licensing Guide</i>
Release notes	<i>Cisco Nexus 7000 Series NX-OS Release Notes</i>

Standards

Standards	Title
IEEE 802.1Q-2006 (formerly known as IEEE 802.1s), IEEE 802.1D-2004 (formerly known as IEEE 802.1w), IEEE 802.1D, IEEE 802.1t	—

MIBs

MIBs	MIBs Link
<ul style="list-style-type: none"> • CISCO-STP-EXTENSION-MIB • BRIDGE-MIB 	To locate and download MIBs, go to the following URL: https://cfnng.cisco.com/mibs .

Feature History for Configuring Rapid PVST+—CLI Version

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

Table 6: Feature History for Configuring Rapid PVST+

Feature Name	Releases	Feature Information
No change	4.2(1)	--
No change	4.1(2)	--