31 Days Before Your CCNA Exam (200-301)

A Day-By-Day Review Guide for the CCNA 200-301 Certification Exam

ciscopress.com

Allan Johnson

CCNA Countdown Calendar

The lines after the countdown number allow you to add the actual calendar days for reference.

31 Networking Models, Devices, and Components	**30** Ethernet Switching	**29** Switch Configuration Basics	**28** IPv4 Addressing
24 EtherChannel and HSRP	**23** DHCP and DNS	**22** Wireless Concepts	**21** WLAN Configuration
17 The Routing Table	**16** Inter-VLAN Routing	**15** Static and Default Route Configuration	**14** OSPF Operation
10 ACL Concepts	**9** ACL Implementation	**8** NAT	**7** WAN, VPN, and IPsec
3 Cloud, Virtualization, and SDN	**2** SDA and Cisco DNA Center	**1** Network Automation	**EXAM** Time _____ Location _____

The BID includes the following fields:

- **Bridge Priority:** A 4-bit field is still used to carry bridge priority. However, the priority is conveyed in discrete values in increments of 4096 instead of discrete values in increments of 1 because only the first 4 most-significant bits are available from the 16-bit field.
- **Extended System ID:** A 12-bit field carrying the VID for PVST+.
- **MAC Address:** A 6-byte field with the MAC address of a single switch.

Rapid PVST+ Operation

In Rapid PVST+, a single instance of RSTP runs for each VLAN. This is why Rapid PVST+ has a very high demand for switch resources (CPU cycles and RAM).

NOTE: Rapid PVST+ is simply the Cisco implementation of RSTP on a per-VLAN basis. The rest of this review uses the terms *RSTP* and *Rapid PVST+* interchangeably.

With RSTP, the IEEE improved the convergence performance of STP from 50 seconds to less than 10 seconds with its definition of Rapid STP (RSTP) in the standard 802.1w. RSTP is identical to STP in the following ways:

- It elects the root switch by using the same parameters and tiebreakers.
- It elects the root port on nonroot switches by using the same rules.
- It elects designated ports on each LAN segment by using the same rules.
- It places each port in either forwarding or discarding state, although RSTP calls the blocking state the discarding state.

RSTP Interface Behavior

The main changes with RSTP can be seen when changes occur in the network. RSTP acts differently on some interfaces based on what is connected to the interface:

- **Edge-type behavior and PortFast:** RSTP improves convergence for edge-type connections by immediately placing the port in forwarding state when the link is physically active.
- **Link-type shared:** RSTP does not do anything differently from STP on link-type shared links. However, because most links between switches today are full duplex, point-to-point, and not shared, this does not matter.
- **Link-type point-to-point:** RSTP improves convergence over full-duplex links between switches. RSTP recognizes the loss of the path to the root bridge through the root port in 6 seconds (based on three times the hello timer value of 2 seconds). RSTP thus recognizes a lost path to the root much more quickly.

RSTP uses different terminology to describe port states. Table 25-5 lists the port states for RSTP and STP.

Table 25-5 RSTP and STP Port States

Operational State	STP State (802.1D)	RSTP State (802.1w)	Forwards Data Frames in This State?
Enabled	Blocking	Discarding	No
Enabled	Listening	Discarding	No
Enabled	Learning	Learning	No
Enabled	Forwarding	Forwarding	Yes
Disabled	Disabled	Discarding	No

RSTP removes the need for listening state and reduces the time required for learning state by actively discovering the network's new state. STP passively waits on new BPDUs and reacts to them during the listening and learning states. With RSTP, the switches negotiate with neighboring switches by sending RSTP messages. The messages enable the switches to quickly determine whether an interface can be immediately transitioned to a forwarding state. In many cases, the process takes only a second or two for the entire RSTP domain.

RSTP Port Roles

RSTP adds three more port roles in addition to the root port and designated port roles defined in STP. Table 25-6 lists and defines the port roles.

Table 25-6 RSTP and STP Port Roles

RSTP Role	STP Role	Definition
Root port	Root port	A single port on each nonroot switch in which the switch hears the best BPDU out of all the received BPDUs
Designated port	Designated port	Of all switch ports on all switches attached to the same segment/collision domain, the port that advertises the "best" BPDU
Alternate port	—	A port on a switch that receives a suboptimal BPDU
Backup port	—	A nondesignated port on a switch that is attached to the same segment/collision domain as another port on the same switch
Disabled	—	A port that is administratively disabled or that is not capable of working for other reasons

Figure 25-6 shows an example of these RSTP port roles.

Figure 25-6 RSTP Port Roles

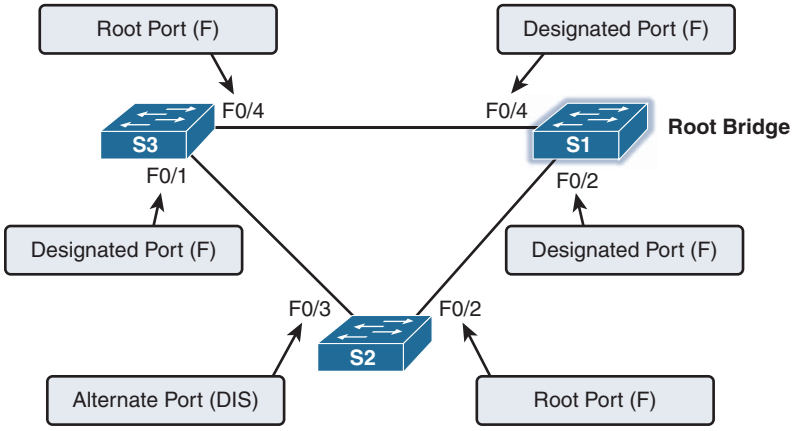

Edge Ports

In addition to the port roles just described, RSTP uses an edge port concept that corresponds to the PVST+ PortFast feature. An edge port connects directly to an end device. Therefore, the switch assumes that no other switch is connected to it. RSTP edge ports should immediately transition to the forwarding state, thereby skipping the time-consuming original 802.1D listening and learning port states. The only caveat is that the port must be a point-to-point link. If it is a shared link, the port is not an edge port, and PortFast should not be configured. Why? Another switch could be added to a shared link—on purpose or inadvertently. Figure 25-7 shows examples of edge ports.

Figure 25-7 Edge Ports in RSTP

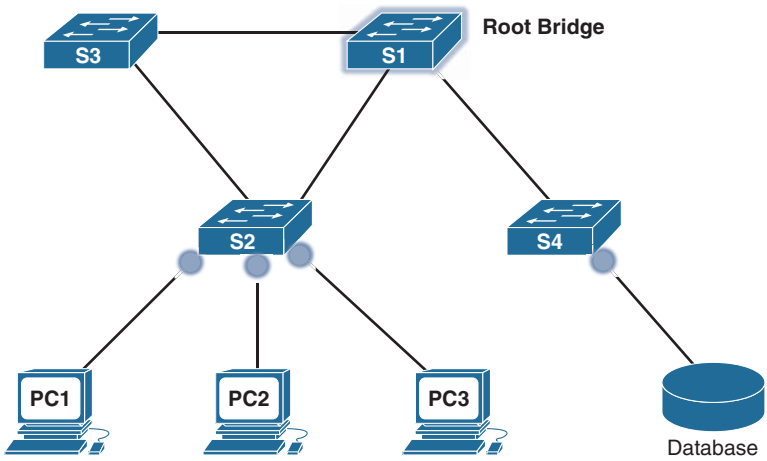

Configuring and Verifying Varieties of STP

By default, all Cisco switches use STP without any configuration by the network administrator. However, because STP runs on a per-VLAN basis, you can take advantage of several options to load balance traffic across redundant links.

STP Configuration Overview

Before you configure or alter the behavior of STP, it is important to know the current default settings listed in Table 25-7.

Table 25-7 Default STP Configuration on the Cisco Catalyst 2960

Feature	Default Setting
Enable state	Enables STP on VLAN 1
Spanning tree mode	PVST+ (Rapid PVST+ and MSTP disabled)
Switch priority	32768
Spanning tree port priority (configurable on a per-interface basis)	128
Spanning tree port cost (configurable on a per-interface basis)	1000 Mbps: 4 100 Mbps: 19 10 Mbps: 100
Spanning tree VLAN port priority (configurable on a per-VLAN basis)	128
Spanning tree VLAN port cost (configurable on a per-VLAN basis)	1000 Mbps: 4 100 Mbps: 19 10 Mbps: 100
Spanning tree timers	Hello time: 2 seconds Forward-delay time: 15 seconds Maximum-aging time: 20 seconds Transmit hold count: 6 BPDUs

Configuring and Verifying the BID

Regardless of which PVST you use, two main configuration options can help you achieve load balancing: the bridge ID and the port cost manipulation. The bridge ID influences the choice of root switch and can be configured per VLAN. Each interface's (per-VLAN) STP cost to reach the root influences the choice of designated port on each LAN segment. Because PVST requires that a separate instance of spanning tree run for each VLAN, the BID field is required to carry VLAN ID (VID) information. This is accomplished by reusing a portion of the Priority field as the extended system ID to carry a VID.

To change the bridge ID, use one of the following commands:

```
Switch(config)# spanning-tree vlan vlan-id root {primary | secondary}
Switch(config)# spanning-tree vlan vlan-id priority priority
```

To change the interface cost, use the following command:

```
Switch(config-if)# spanning-tree vlan vlan-id cost cost
```

Figure 25-8 shows a simple three-switch STP topology without redundant links.

Figure 25-8 STP Topology

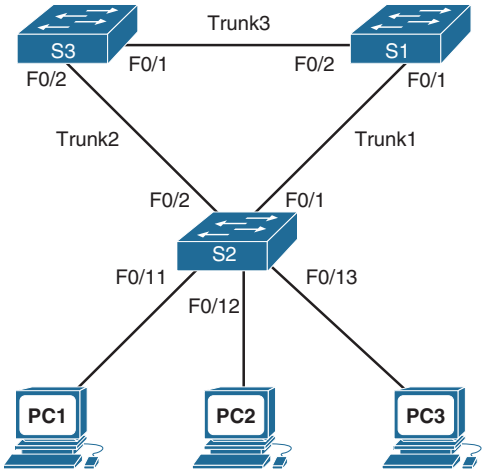

The network administrator wants to ensure that S1 is always the root bridge and S2 is the backup root bridge. The following commands achieve this objective:

```
S1(config)# spanning-tree vlan 1 root primary
!---------
S2(config)# spanning-tree vlan 1 root secondary
```

The **primary** keyword automatically sets the priority to 24576 or to the next 4096 increment value below the lowest bridge priority detected on the network.

The **secondary** keyword automatically sets the priority to 28672, assuming that the rest of the network is set to the default priority of 32768.

Alternatively, the network administrator can explicitly configure the priority value in increments of 4096 between 0 and 65536 using the following command:

```
S1(config)# spanning-tree vlan 1 priority 24576
!---------
S2(config)# spanning-tree vlan 1 priority 28672
```

NOTE: In this example, these commands changed the priority values only for VLAN 1. Additional commands must be entered for each VLAN to take advantage of load balancing.

To verify the current spanning tree instances and root bridges, use the **show spanning-tree** command (see Example 25-1).

Example 25-1 Verifying Spanning Tree Configurations

```
S1# show spanning-tree

VLAN0001
  Spanning tree enabled protocol ieee
  Root ID    Priority    24577
             Address     001b.5302.4e80
             This bridge is the root
             Hello Time  2 sec Max Age 20 sec Forward Delay 15 sec

  Bridge ID Priority    24577 (priority 24576 sys-id-ext 1)
             Address     001b.5302.4e80
             Hello Time  2 sec Max Age 20 sec Forward Delay 15 sec
             Aging Time  300

Interface         Role Sts Cost       Prio.NbrType
----------------- ---- --- ---------- -------- --------------------------------
Fa0/1             Desg  FWD 19         128.1 P2p
Fa0/2             Desg  FWD 19         128.2 P2p
```

Because an extended system ID is used in the BID, the value of the priority includes the addition of the VLAN ID. Therefore, a priority of 24576 plus a VLAN of 1 results in a priority output of 24577.

Configuring PortFast and BPDU Guard

To speed convergence for access ports when they become active, you can use Cisco's proprietary PortFast technology. After PortFast is configured and a port is activated, the port immediately transitions from the blocking state to the forwarding state.

In a valid PortFast configuration, BPDUs should never be received because receipt of a BPDU indicates that another bridge or switch is connected to the port, potentially causing a spanning tree loop. When it is enabled, BPDU Guard puts the port in an errdisabled (error-disabled) state upon receipt of a BPDU. This effectively shuts down the port. The BPDU Guard feature provides a secure response to invalid configurations because you must manually put the interface back into service.

Example 25-2 shows the interface commands to configure PortFast and BPDU Guard on S2 in Figure 25-8.