CHAPTER 12

Advanced Spanning Tree Protocol

Familiarity with the IEEE 802.1D STP standard is essential because that protocol is used universally to maintain loop-free bridged and switched networks. However, it is now considered a legacy protocol, offering topology change and convergence times that are not as acceptable as they once were.

This chapter discusses the many STP enhancements that are available in new standards. The Rapid STP (RSTP) is presented first because it provides the foundation for efficient STP activity. The Multiple STP (MST or MSTP) is also discussed here. MST allows VLANs to be individually mapped into STP instances, while RSTP operates in the background. These two protocols allow a Layer 2 campus network to undergo change quickly and efficiently, with little downtime for today's applications.

"Do I Know This Already?" Quiz

The purpose of the "Do I Know This Already?" quiz is to help you decide if you need to read the entire chapter. If you already intend to read the entire chapter, you do not necessarily need to answer these questions now.

The quiz, derived from the major sections in the "Foundation Topics" portion of the chapter, helps you determine how to spend your limited study time.

Table 12-1 outlines the major topics discussed in this chapter and the "Do I Know This Already?" quiz questions that correspond to those topics.

Table 12-1 *"Do I Know This Already?" Foundation Topics Section-to-Question Mapping*

Foundation Topics Section	Questions Covered in This Section
Rapid STP	1–8
Multiple STP (MST)	9–12

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark this question wrong. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might give you a false sense of security.

1. Which one of the following enables the use of RSTP?
   a. PVST+
   b. 802.1D
   c. CST
   d. MST

2. Upon which standard is RSTP based?
   a. 802.1Q
   b. 802.1D
   c. 802.1w
   d. 802.1s

3. Which of the following is not a port state in RSTP?
   a. Listening
   b. Learning
   c. Discarding
   d. Forwarding

4. When a switch running RSTP receives an 802.1D BPDU, what happens?
   a. The BPDU is discarded or dropped.
   b. An ICMP message is returned.
   c. The switch begins to use 802.1D rules on that port.
   d. The switch disables RSTP.

5. When does an RSTP switch consider a neighbor to be down?

   a. After three BPDUs are missed
   b. After six BPDUs are missed
   c. After the Max Age timer expires
   d. After the Forward Timer expires

6. Which process is used during RSTP convergence?

   a. BPDU propagation
   b. Synchronization
   c. Forward Timer expiration
   d. BPDU acknowledgments

7. What causes RSTP to view a port as a point-to-point port?

   a. Port speed
   b. Port media
   c. Port duplex
   d. Port priority

8. Which of the following events triggers a topology change with RSTP on a nonedge port?

   a. A port comes up or goes down.
   b. A port comes up.
   c. A port goes down.
   d. A port moves to the forwarding state.

9. Which of the following is not a characteristic of MST?

   a. A reduced number of STP instances
   b. Fast STP convergence
   c. Eliminates the need for CST
   d. Interoperability with PVST+

10. Which of the following standards defines the MST protocol?
    a. 802.1Q
    b. 802.1D
    c. 802.1w
    d. 802.1s

11. How many instances of STP are supported in the Cisco implementation of MST?
    a. 1
    b. 16
    c. 256
    d. 4096

12. What switch command can be used to change from PVST+ to MST?
    a. **spanning-tree mst enable**
    b. **no spanning-tree pvst+**
    c. **spanning-tree mode mst**
    d. **spanning-tree mst**

You can find the answers to the "Do I Know This Already?" quiz in Appendix A, "Answers to Chapter 'Do I Know This Already?' Quizzes and Q&A Sections." The suggested choices for your next step are as follows:

- **10 or less overall score**—Read the entire chapter. This includes the "Foundation Topics," "Foundation Summary," and "Q&A" sections.
- **11 or 12 overall score**—If you want more review on these topics, skip to the "Foundation Summary" section and then go to the "Q&A" section at the end of the chapter. Otherwise, move to Chapter 13, "Multilayer Switching."

Foundation Topics

Rapid Spanning Tree Protocol (RSTP)

The IEEE 802.1D Spanning Tree Protocol was designed to keep a switched or bridged network loop free, with adjustments made to the network topology dynamically. A topology change typically takes 30 seconds, where a port moves from the Blocking state to the Forwarding state after two intervals of the Forward Delay timer. As technology has improved, 30 seconds has become an unbearable length of time to wait for a production network to failover or "heal" itself during a problem.

The IEEE 802.1w standard was developed to take 802.1D's principle concepts and make the resulting convergence much faster. This is also known as the Rapid Spanning Tree Protocol (RSTP). RSTP defines how switches must interact with each other to keep the network topology loop free, in a very efficient manner. Like 802.1D, RSTP's basic functionality can be applied as a single or multiple instances. This can be done as the IEEE 802.1s Multiple Spanning Tree (MST), covered in this chapter, and also as the Cisco-proprietary, Rapid Per-VLAN Spanning Tree Protocol (RPVST+). RSTP operates consistently in each, but replicating RSTP as multiple instances requires different approaches.

RSTP Port Behavior

In 802.1D, each switch port is assigned a role and a state at any given time. Depending on the port's proximity to the Root Bridge, it takes on one of the following roles:

- Root Port
- Designated Port
- Blocking Port (neither Root nor Designated).

The Cisco-proprietary UplinkFast feature also reserved a hidden Alternate Port role for ports that offered parallel paths to the Root but were in the Blocking state.

Recall that each switch port is also assigned one of five possible states:

- Disabled
- Blocking
- Listening
- Learning
- Forwarding

Only the Forwarding state allows data to be sent and received. A port's state is somewhat tied to its role. For example, a Blocking Port cannot be a Root Port or a Designated Port.

RSTP achieves its rapid nature by letting each switch interact with its neighbors through each port. This interaction is performed based on a port's role, not strictly on the BPDUs that are relayed from the Root Bridge. After the role is determined, each port can be given a state that determines what it does with incoming data.

The Root Bridge in a network using RSTP is elected just as with 802.1D—by the lowest Bridge ID. After all switches agree on the identity of the Root, the following port roles are determined:

- **Root Port**—The one switch port on each switch that has the best root path cost to the Root. This is identical to 802.1D. (By definition, the Root Bridge has no Root Ports.)
- **Designated Port**—The switch port on a network segment that has the best root path cost to the Root.
- **Alternate Port**—A port that has an alternate path to the Root, different than the path the Root Port takes. This path is less desirable than that of the Root Port. (An example of this is an access layer switch with two uplink ports; one becomes the Root Port, the other is an Alternate Port.)
- **Backup Port**—A port that provides a redundant (but less desirable) connection to a segment where another switch port already connects. If that common segment is lost, the switch might or might not have a path back to the Root.

RSTP defines port states only according to what the port does with incoming frames. (Naturally, if incoming frames are ignored or dropped, so are outgoing frames.) Any port role can have any of these port states:

- **Discarding**—Incoming frames are simply dropped; no MAC addresses are learned. (This state combines the 802.1D Disabled, Blocking, and Listening states, as all three did not effectively forward anything. The Listening state is not needed, because RSTP can quickly negotiate a state change without listening for BPDUs first.)
- **Learning**—Incoming frames are dropped, but MAC addresses are learned.
- **Forwarding**—Incoming frames are forwarded according to MAC addresses that have been (and are being) learned.

BPDUs in RSTP

In 802.1D, BPDUs basically originate from the Root Bridge and are relayed by all switches down through the tree. It is because of this propagation of BPDUs that 802.1D convergence must wait for steady-state conditions before proceeding.

RSTP uses the 802.1D BPDU format for backward-compatibility. However, some previously unused bits in the Message Type field are used. The sending switch port identifies itself by its RSTP role and state. The BPDU version is also set to 2, to distinguish RSTP BPDUs from 802.1D BPDUs. Also, RSTP uses an interactive process so that two neighboring switches can negotiate state changes. Some BPDU bits are used to flag messages during this negotiation.

BPDUs are sent out every switch port at Hello Time intervals, regardless of whether BPDUs are received from the Root. In this way, any switch anywhere in the network can play an active role in maintaining the topology. Switches can also expect to receive regular BPDUs from their neighbors. When three BPDUs are missed in a row, that neighbor is presumed to be down, and all information related to the port leading to the neighbor is immediately aged out. This means that a switch can detect a neighbor failure in three Hello intervals (default 6 seconds), versus the Max Age Timer interval (default 20 seconds) for 802.1D.

Because RSTP distinguishes its BPDUs from 802.1D BPDUs, it can coexist with switches still using 802.1D. Each port attempts to operate according to the STP BPDU that is received. For example, when an 802.1D BPDU (version 0) is received on a port, that port begins to operate according to the 802.1D rules. However, each port has a measure that locks the protocol in use for the duration of the migration delay timer. This keeps the protocol type from flapping or toggling during a protocol migration. After the timer expires, the port is free to change protocols if needed.

RSTP Convergence

The convergence of STP in a network is the process that takes all switches from a state of independence (each thinks it must be the STP Root) to one of uniformity, where each switch has a place in a loop-free tree topology. You can think of convergence as a two-stage process:

1. One common Root Bridge must be "elected," and all switches must know about it.
2. The state of every switch port in the STP domain must be brought from a Blocking state to the appropriate state to prevent loops.

Convergence generally takes time, as messages are propagated from switch to switch. The traditional 802.1D STP also requires the expiration of several timers before switch ports can be safely allowed to forward data.

RSTP takes a different approach when a switch needs to decide how to participate in the tree topology. When a switch first joins the topology (perhaps it was just powered up) or has detected a failure in the existing topology, RSTP requires it to base its forwarding decisions on the type of port.

Port Types

Every switch port can be considered one of the following types:

- **Edge Port**—A port at the "edge" of the network, where only a single host connects. Traditionally, this has been identified by enabling the STP PortFast feature. RSTP keeps the PortFast concept for familiarity. By definition, the port cannot form a loop as it connects to one host, so it can be immediately placed in the Forwarding state. However, if a BPDU is ever received on an edge port, the port immediately loses its edge port status.
- **Root Port**—The port that has the best cost to the root of the STP instance. Only one Root Port can be selected and active at any time, although alternate paths to the root can exist through other ports. If alternate paths are detected, those ports are identified as Alternate Root Ports and can be immediately placed in the Forwarding state when the existing Root Port fails.
- **Point-to-Point Port**—Any port that connects to another switch and becomes a Designated Port. A quick handshake with the neighboring switch, rather than a timer expiration, decides the port state. BPDUs are exchanged back and forth in the form of a proposal and an agreement. One switch proposes that its port becomes a Designated Port; if the other switch agrees, it replies with an agreement message.

  Point-to-point ports are automatically determined by the duplex mode in use. Full-duplex ports are considered point-to-point because only two switches can be present on the link. STP convergence can quickly occur over a point-to-point link through RSTP handshake messages.

  Half-duplex ports, on the other hand, are considered to be on a shared media with possibly more than two switches present. They are not point-to-point ports. STP convergence on a half-duplex port must occur between several directly connected switches. Therefore, the traditional 802.1D style convergence must be used. This results in a slower response because the shared-media ports must go through the fixed listening and learning state time periods.

It's easy to see how two switches can quickly converge to a common idea of which one is the Root and which one will have the Designated Port after just a single exchange of BPDUs. What about a larger network, where 802.1D BPDUs would normally have to be relayed from switch to switch?

RSTP handles the complete STP convergence of the network as a propagation of handshakes over point-to-point links. When a switch needs to make an STP decision, a handshake is made with the nearest neighbor. After that is successful, the handshake sequence is moved to the next switch and the next, as an ever-expanding wave moving toward the network's edges.

During each handshake sequence, a switch must take measures to be completely sure it will not introduce a bridging loop before moving the handshake out. This is done through a synchronization process.

Synchronization

To participate in RSTP convergence, a switch must decide the state of each of its ports. Nonedge ports begin in the Discarding state. After BPDUs are exchanged between the switch and its neighbor, the Root Bridge can be identified. If a port receives a superior BPDU from a neighbor, that port becomes the Root Port.

For each nonedge port, the switch exchanges a proposal-agreement handshake to decide the state of each end of the link. Each switch assumes that its port should become the Designated Port for the segment, and a proposal message (a Configuration BPDU) is sent to the neighbor suggesting this.

When a switch receives a proposal message on a port, the following sequence of events occurs (Figure 12-1 shows the sequence, based around the center Catalyst switch):

1. If the proposal's sender has a superior BPDU, the local switch realizes that the sender should be the Designated Switch (having the Designated Port), and that its own port must become the new Root Port.
2. Before the switch agrees to anything, it must first synchronize itself with the topology.
3. All nonedge ports are immediately moved into the Discarding (blocking) state so that no bridging loops can form.
4. An agreement message (a Configuration BPDU) is sent back to the sender, indicating that the switch is in agreement with the new Designated Port choice. This also tells the sender that the switch is in the process of synchronizing itself.
5. The Root Port is immediately moved to the Forwarding state. The sender's port can also immediately begin forwarding.
6. For each nonedge port that is currently in the Discarding state, a proposal message is sent to the respective neighbor.
7. An agreement message is expected and received from a neighbor on a nonedge port.
8. The nonedge port is immediately moved to the Forwarding state.

Notice how the RSTP convergence begins with a switch sending a proposal message. The recipient of the proposal must synchronize itself by effectively isolating itself from the rest of the topology. All nonedge ports are blocked until a proposal message can be sent, causing the nearest neighbors to synchronize themselves. This creates a moving "wave" of synchronizing switches, which can quickly decide to start forwarding on their links only if their neighbors agree. Figure 12-2 shows how the synchronization wave travels through a network at three successive time intervals. Isolating the switches along the traveling wave inherently prevents bridging loops.

Figure 12-1 *Sequence of Events During RSTP Convergence*

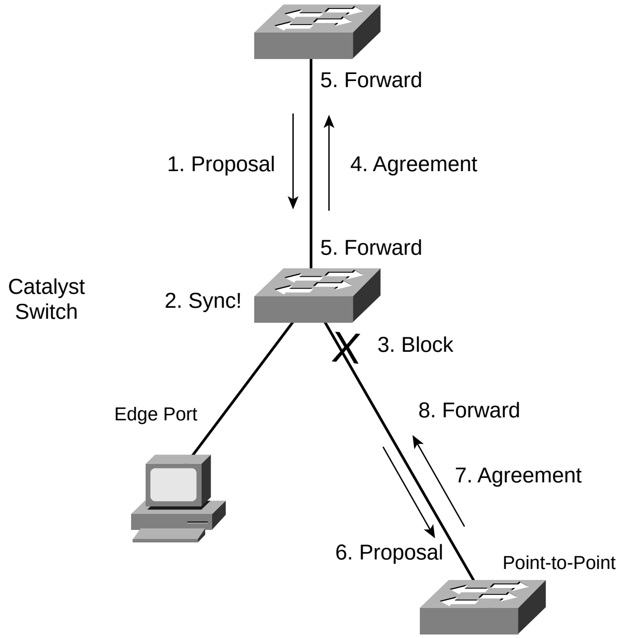

The entire convergence process happens quickly, at the speed of BPDU transmission, without the use of any timers. A Designated Port that sends a proposal message might not receive an agreement message reply. Suppose the neighboring switch does not understand RSTP or has a problem replying. The sending switch must then become overly cautious and begin playing by the 802.1D rules—the port must be moved through the legacy Listening and Learning states (using the Forward Delay timer) before moving to the Forwarding state.

Topology Changes and RSTP

Recall that when an 802.1D switch detects a port state change (either up or down), it signals the Root Bridge by sending topology change notification (TCN) BPDUs. The Root Bridge must then signal a topology change by sending out a TCN message that is relayed to all switches in the STP domain.

RSTP detects a topology change only when a nonedge port transitions to the Forwarding state. This might seem odd because a link failure is not used as a trigger. RSTP uses all of its rapid convergence mechanisms to prevent bridging loops from forming. Therefore, topology changes are detected only so that bridging tables can be updated and corrected as hosts appear first on a failed port and then on a different functioning port.

Figure 12-2 *RSTP Synchronization Traveling Through a Network*

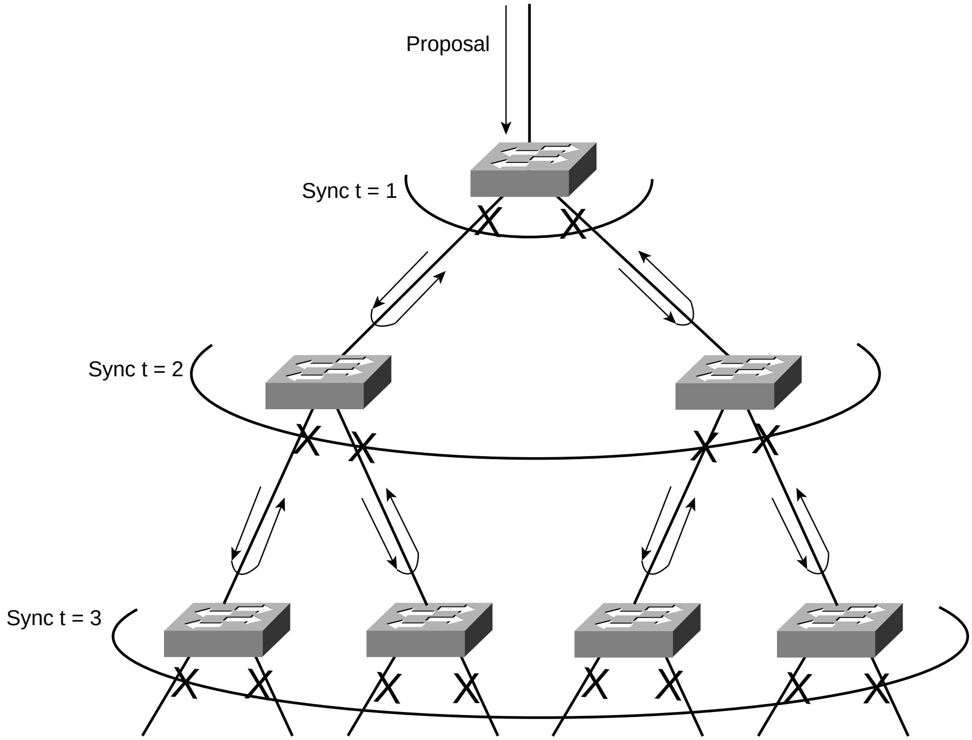

When a topology change is detected, a switch must propagate news of the change to other switches in the network so they can correct their bridging tables, too. This process is similar to the convergence and synchronization mechanism—topology change (TC) messages propagate through the network in an ever-expanding wave.

BPDUs, with their TC bit set, are sent out all of the nonedge designated ports. This is done until the "TC While" timer expires, after two times the Hello time. This notifies neighboring switches of the new link and the topology change. In addition, all MAC addresses associated with the nonedge Designated Ports are flushed from the content-addressable memory (CAM) table. This forces the addresses to be relearned after the change, in case hosts now appear on a different link.

All neighboring switches that receive the TC messages must also flush the MAC addresses learned on all ports except the one that received the TC message. Those switches must then send TC messages out their nonedge Designated Ports, and so on.

RSTP Configuration

By default, a switch operates in the Per VLAN Spanning Tree Plus (PVST+) mode using traditional 802.1D STP. Therefore, RSTP cannot be used until a different Spanning Tree mode (MST or RPVST+) is enabled. Remember that RSTP is just the underlying mechanism that a Spanning Tree mode can use to detect topology changes and converge a network into a loop-free topology.

The only configuration changes related to RSTP affect the port or link type. The link type is used to determine how a switch negotiates topology information with its neighbors.

To configure a port as an RSTP edge port, use the following interface configuration command:

```
Switch(config-if)# spanning-tree portfast
```

You should already be familiar with this command from the 802.1D STP configuration. After PortFast is enabled, the port is considered to have only one host and is positioned at the edge of the network.

By default, RSTP automatically decides that a port is a point-to-point link if it is operating in full-duplex mode. Ports connecting to other switches are usually full-duplex because there are only two switches on the link. However, you can override the automatic determination if needed. For example, a port connecting to one other switch might be operating at half-duplex for some reason. To force the port to act as a point-to-point link, use the following interface configuration command:

```
Switch(config-if)# spanning-tree link-type point-to-point
```

Multiple Spanning Tree (MST) Protocol

Chapter 9 covered two “flavors” of Spanning Tree implementations—IEEE 802.1Q and PVST+—both based on the 802.1D STP. These also represent the two extremes of Spanning Tree Protocol operation in a network:

- **802.1Q**—Only a single instance of STP is used for all VLANs. If there are 500 VLANs, only one instance of STP will be running. This is called the Common Spanning Tree (CST) and operates over the trunk’s native VLAN.
- **PVST+**—One instance of STP is used for each active VLAN in the network. If there are 500 VLANs, 500 independent instances of STP will be running.

In most networks, each switch has a redundant path to another switch. For example, an access layer switch usually has two uplinks, each connecting to a different distribution or core layer switch. If 802.1Q’s CST is used, only one STP instance will run. That means there is only one loop-free topology at any given time, and that only one of the two uplinks in the access layer switch will be forwarding. The other uplink will always be blocking.

Obviously, arranging the network so that both uplinks can be used simultaneously would be best. One uplink should carry one set of VLANs, while the other carries a different set, as a type of load balancing.

PVST+ seems more attractive to meet that goal because it allows different VLANs to have different topologies, so that each uplink can be forwarding. But, think of the consequences—as the number of VLANs increases, so does the number of independent STP instances. Each instance uses some amount of the switch CPU and memory resources. The more instances in use, the less CPU resources available for switching.

Beyond that, what is the real benefit of having 500 STP topologies for 500 VLANs, when only a small number of possible topologies exist for a switch with two uplinks? Figure 12-3 shows a typical network with an access layer switch connecting to a pair of core switches. Two VLANs are in use, with the Root Bridges configured to support load balancing across the two uplinks. The right portion of the figure shows every possible topology for VLANs A and B. Notice that because the access layer switch has only two uplinks, only two topologies actually matter—one where the left uplink forwards, and one where the right uplink forwards.

Figure 12-3 *The Possible STP Topologies for Two VLANs*

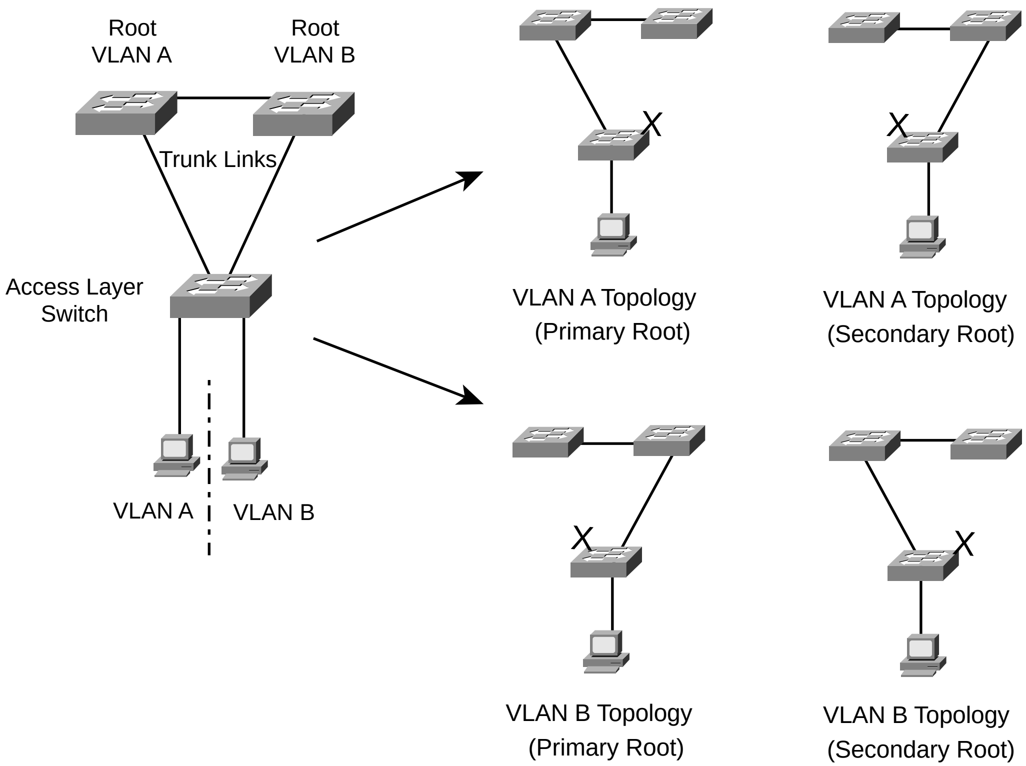

Notice also that the number of useful topologies is independent of the number of VLANs. If 10 or 100 VLANs were used in the figure, there would still be only two possible outcomes at the access layer switch. Therefore, running 10 or 100 instances of STP when only a couple would suffice is rather wasteful.

The Multiple Spanning Tree Protocol (MST or MSTP) was developed to address the lack of and surplus of STP instances. As a result, the network administrator can configure exactly the number of STP instances that make sense for the enterprise network—no matter how many VLANs are in use. MST is defined in the IEEE 802.1s standard.

MST Overview

MST is built on the concept of mapping one or more VLANs to a single STP instance. Multiple instances of STP can be used (hence the name MST), with each instance supporting a different group of VLANs.

For the network shown in Figure 12-3, only two MST instances would be needed. Each could be tuned to result in a different topology, so that Instance 1 would forward on the left uplink, while Instance 2 would forward on the right uplink. Therefore, VLAN A would be mapped to Instance 1, and VLAN B to Instance 2.

To implement MST in a network, you need to determine the following:

- The number of STP instances needed to support the desired topologies.
- Whether to map a set of VLANs to each instance.

MST Regions

MST is different than 802.1Q and PVST+, although it can interoperate with them. If a switch is configured to use MST, it must somehow figure out which of its neighbors are using which type of STP. This is done by configuring switches into common MST regions, where every switch in a region runs MST with compatible parameters.

In most networks, a single MST region is sufficient, although you can configure more than one region. Within the region, all switches must run the instance of MST that is defined by the following attributes:

- MST configuration name (32 characters)
- MST configuration revision number (0 to 65535)
- MST instance-to-VLAN mapping table (4096 entries)

If two switches have the same set of attributes, they belong to the same MST region. If not, they belong to two independent regions.

MST BPDUs contain configuration attributes so that switches receiving BPDUs can compare them against their local MST configurations. If the attributes match, the STP instances within MST can be shared as part of the same region. If not, a switch is seen to be at the MST region boundary, where one region meets another or one region meets traditional 802.1D STP.

NOTE The entire MST instance-to-VLAN mapping table is not sent along in the BPDUs because the instance mappings must be configured on each switch. Instead, a digest, or a code computed from the table contents, is sent. As the contents of the table change, the digest value will be different. Therefore, a switch can quickly compare a received digest to its own to see if the advertised table is the same or different.

Spanning Tree Instances Within MST

MST was designed to interoperate with all other forms of STP. Therefore, it must also support STP instances from each. This is where MST can get confusing. Think of the entire enterprise network having a single CST topology, such that one instance of STP represents any and all VLANs and MST regions present. The CST serves to maintain a common loop-free topology, while integrating all forms of STP that might be in use.

To do this, CST must regard each MST region as a single "black box" bridge because it has no idea what is inside the region, nor does it care. CST only maintains a loop-free topology with the links that connect the regions to each other and to standalone switches running 802.1Q CST.

IST Instances

Something other than CST must work out a loop-free topology inside each MST region. Within a single MST region, an Internal Spanning Tree (IST) instance runs to work out a loop-free topology between the links where CST meets the region boundary and all switches inside the region. Think of the IST instance as a locally significant CST, bounded by the edges of the region.

The IST presents the entire region as a single virtual bridge to the CST outside. BPDUs are exchanged at the region boundary only over the native VLAN of trunks, as if a single CST were in operation. And, indeed, it is.

Figure 12-4 shows the basic concept behind the IST instance. The network at the left has an MST region, where several switches are running compatible MST configurations. Another switch is outside the region because it is running only the CST from 802.1Q.

The same network is shown at the right, where the IST has produced a loop-free topology for the network inside the region. The IST makes the internal network look like a single bridge (the "big switch" in the cloud) that can interface with the CST running outside the region.

Figure 12-4 *Concepts Behind the IST Instance*

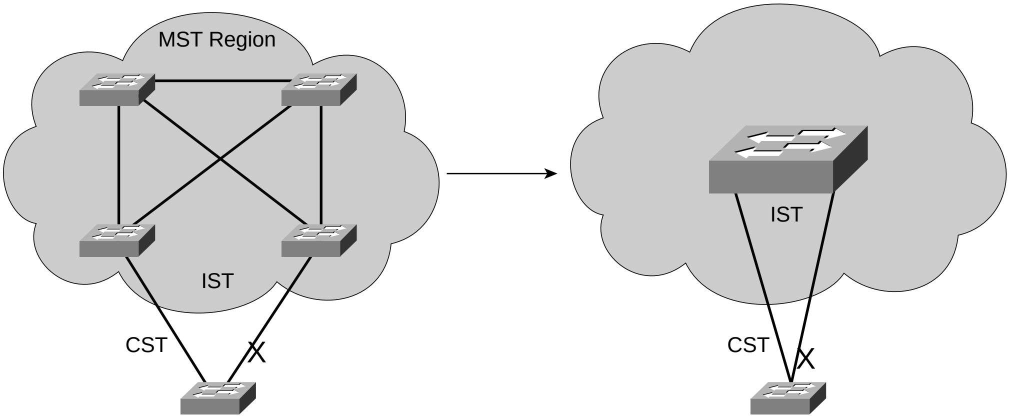

MST Instances

Recall that the whole idea behind MST is the capability to map multiple VLANs to a smaller number of STP instances. Inside a region, the actual MST instances (MSTIs) exist alongside the IST. Cisco supports a maximum of 16 MSTIs in each region. IST always exists as MSTI number 0, leaving MSTI 1 through 15 available for use.

Figure 12-5 shows how different MSTIs can exist within a single MST region. The left portion of the figure is identical to that of Figure 12-4. In this network, two MST instances, MSTI 1 and MSTI 2, are configured with different VLANs mapped to each. Their topologies follow the same structure as the network on the left side of the figure, but each has converged differently. Notice that within the MST cloud, we now have three independent STP instances coexisting—MSTI1, MSTI 2, and the IST.

Only the IST (MSTI 0) is allowed to send and receive MST BPDUs. Information about each of the other MSTIs is appended to the MST BPDU as an M-record. Therefore, even if a region has all 16 instances active, only one BPDU is needed to convey STP information about them all.

Each of the MSTIs are significant only within a region, even if an adjacent region has the same MSTIs in use. In other words, the MSTIs combine with the IST only at the region boundary to form a subtree of the CST. That means only IST BPDUs are sent into and out of a region.

Figure 12-5 *Concepts Behind MST Instances*

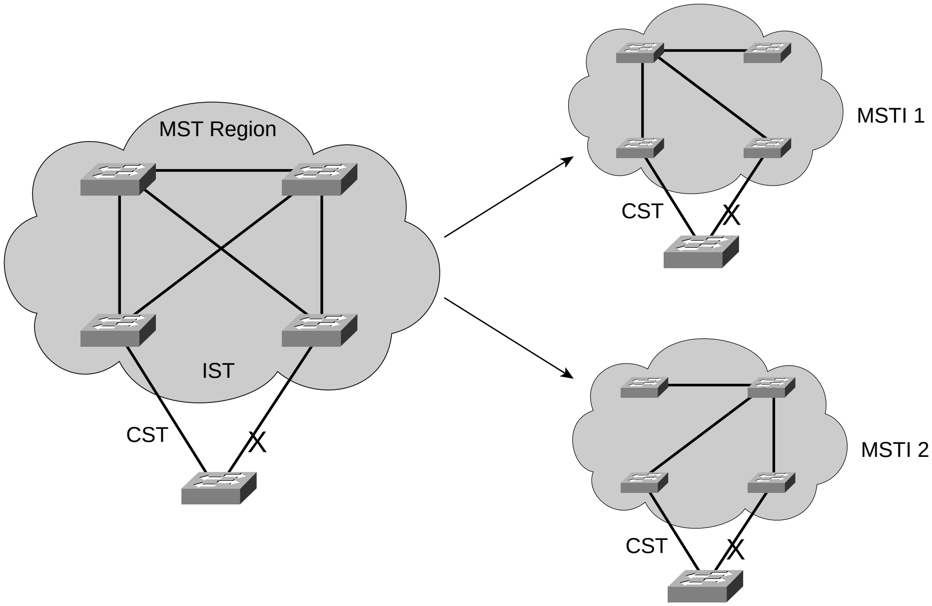

What if an MST region connects with a switch running PVST+? MST can detect this situation by listening to the received BPDUs. If BPDUs are heard from more than one VLAN (the CST), PVST+ must be in use. When the MST region sends a BPDU toward the PVST+ switch, the IST BPDUs are replicated into all of the VLANs on the PVST+ switch trunk.

> **NOTE** Keep in mind that the IST instance is active on *every* port on a switch. Even if a port does not carry VLANs that have been mapped to the IST, IST must still be running on the port.
>
> Also, by default, all VLANs are mapped to the IST instance. You must explicitly map them to other instances if needed.

MST Configuration

You must manually configure the MST configuration attributes on each switch in a region. There is currently no method to propagate this information from one switch to another, as is done with a protocol like VLAN Trunking Protocol (VTP). To define the MST region, use the followig configuration commands in order:

Step 1 Enable MST on the switch:

```
Switch(config)# spanning-tree mode mst
```

Step 2 Enter the MST configuration mode:

```
Switch(config)# spanning-tree mst configuration
```

Step 3 Assign a region configuration name (up to 32 characters):

```
Switch(config-mst)# name name
```

Step 4 Assign a region configuration revision number (0 to 65,535):

```
Switch(config-mst)# revision version
```

The configuration revision number gives you a means to track changes to the MST region configuration. Each time you make changes to the configuration, you should increase the number by one. Remember that the region configuration (including the revision number) must match on all switches in the region. Therefore, you also need to update the revision numbers on the other switches to match.

Step 5 Map VLANs to an MST instance:

```
Switch(config-mst)# instance instance-id vlan vlan-list
```

The *instance-id* (0 to 15) carries topology information for the VLANs listed in *vlan-list*. The list can contain one or more VLANs separated by commas. You can also add a range of VLANs to the list by separating numbers with a hyphen. VLAN numbers can range from 1 to 4094. (Remember that by default, all VLANs are mapped to instance 0, the IST.)

Step 6 Show the pending changes you have made:

```
Switch(config-mst)# show pending
```

Step 7 Exit the MST configuration mode; commit the changes to the active MST region configuration:

```
Switch(config-mst)# exit
```

After MST is enabled and configured, PVST+ operation stops and the switch changes to RSTP operation. A switch cannot run both MST and PVST+ at the same time.

You can also tune the parameters that MST uses when it interacts with CST or traditional 802.1D. The parameters and timers are identical to those discussed in Chapter 10, “Spanning Tree Configuration.” In fact, the commands are very similar except for the addition of the **mst** keyword and the *instance-id*. Rather than tuning STP for a VLAN instance, you use an MST instance.

Table 12-2 summarizes the commands as a quick reference. Notice that the timer configurations are applied to MST as a whole, and not to a specific MST instance. This is because all instance timers are defined through the IST instance and BPDUs.

Table 12-2 *MST Configuration Commands*

Task	Command Syntax
Set Root Bridge (macro)	**spanning-tree mst** *instance-id* **root** {**primary** | **secondary**} [**diameter** *diameter*]
Set Bridge Priority	**spanning-tree mst** *instance-id* **priority** *bridge-priority*
Set Port Cost	**spanning-tree mst** *instance-id* **cost** *cost*
Set Port Priority	**spanning-tree mst** *instance-id* **port-priority** *port-priority*
Set STP Timers	**spanning-tree mst hello-time** *seconds* **spanning-tree mst forward-time** *seconds* **spanning-tree mst max-age** *seconds*

Foundation Summary

The Foundation Summary is a collection of information that provides a convenient review of many key concepts in this chapter. If you are already comfortable with the topics in this chapter, this summary can help you recall a few details. If you just read this chapter, this review should help solidify some key facts. If you are doing your final preparation before the exam, this information is a convenient way to review the day before the exam.

RSTP port roles:

- Root Port
- Designated Port
- Alternate Port
- Backup Port

RSTP port states:

- Discarding
- Learning
- Forwarding

Table 12-3 *RSTP Configuration Commands*

Task	Command Syntax
Define an edge port	**spanning-tree portfast**
Override a port type	**spanning-tree link-type point-to-point**

STP instances involved with MST:

- **Common Spanning Tree (CST)**—Used to maintain a single loop-free topology for the entire network
- **Internal Spanning Tree (IST)**—Used like CST to maintain a single, loop-free topology *inside* an MST region
- **MST Instances (MSTIs)**—Used inside an MST region to maintain loop-free topologies for sets of mapped VLANs

Table 12-4 *MST Region Configuration Commands*

Task	Command Syntax
Enable MST on a switch	**spanning-tree mode mst**
Enter MST configuration mode	**spanning-tree mst configuration**
Name the MST region	**name** *name*
Set the configuration revision number	**revision** *version*
Map VLANs to an MST instance	**instance** *instance-id* **vlan** *vlan-list*
Confirm new MST configuration changes	**show pending**
Commit new MST changes	**exit**

Table 12-5 *MST Tuning Configuration Commands*

Task	Command Syntax
Set the Root Bridge	**spanning-tree mode mst** *instance-id* **root** {**primary** | **secondary**} [**diameter** *diameter*]
Set Bridge Priority	**spanning-tree mst** *instance-id* **priority** *bridge-priority*
Set Port Cost	**spanning-tree mst** *instance-id* **cost** *cost*
Set Port Priority	**spanning-tree mst** *instance-id* **port-priority** *port-priority*
Set STP Timers	**spanning-tree mst hello-time** *seconds* **spanning-tree mst forward-time** *seconds* **spanning-tree mst max-age** *seconds*

Chapter 12

"Do I Know This Already?"

1. d
2. c
3. a
4. c
5. a
6. b
7. c
8. d
9. c
10. d
11. b
12. c

Q&A

1. What is synchronization in RSTP?

   As RSTP works to converge a switched network, each switch effectively isolates itself from the next layer of neighbors until an agreement can be reached about who will have the designated port on each segment.

2. What is an Alternate Port?

   A port with an alternate path to the Root. The path is less desirable than the one through the Root Port but is flagged for immediate use if the Root Port path should fail.

3. What is the difference between an Alternate Port and a Backup Port?

   An Alternate Port connects to a different segment than the Root Port to provide an alternate path to the Root. A Backup Port connects to the same segment as another port on the local switch to provide another path out of the switch, but not necessarily another path back to the Root.

4. Can a switch port be a Designated Port and be in the Discarding state?

   Yes; RSTP removes the linkage between a port's role and its state. In 802.1D, a Designated Port must be forwarding, but RSTP doesn't make the same requirement.

5. Which of the following ports can participate in RSTP synchronization?

   a. Root Port
   b. Designated Port
   c. Edge Port
   d. Nonedge Port
   e. Point-to-point Port

   Answer: a, b, d, e

6. What two messages must be exchanged during RSTP synchronization?

   Proposal and Agreement

7. After an agreement message is received from a neighboring switch, how much time elapses before the port can begin forwarding? (Consider any timers that must expire or other conditions that must be met.)

   The port will be moved to the Forwarding state immediately after the agreement message is received. With RSTP, no other conditions are necessary because two switches have completed a quick handshake by exchanging proposal and agreement messages.

8. After a switch receives news of a topology change, how long does it wait to flush entries out of its CAM table?

   The switch flushes entries immediately, rather than employing the timer reduction that 802.1D uses.

9. What command configures a port as an RSTP edge port?

   spanning-tree portfast

10. Suppose interface FastEthernet 0/1 is in half-duplex mode, but you want it to be considered a point-to-point link for RSTP. What command will accomplish this?

    spanning-tree link-type point-to-point

11. Put the following in order of the number of supported STP instances, from lowest to highest:

a. MST

b. PVST+

c. CST

d. 802.1D

Answer; d, c, a, b (d and c both have a single instance.)

12. What three parameters must be configured to uniquely define an MST region?

The region name, configuration revision number, and the instance-to-VLAN mappings

13. What parameter does a switch examine to see if its neighbors have the same VLAN to MST instance mappings? How is that information passed among switches?

The VLAN-to-instance mapping is kept in a table of 4096 entries. This information is passed along in the MST BPDUs. Rather than passing the entire table, switches include only a digest of their current table contents.

14. Which MST instance in a region corresponds to the CST of 802.1Q?

The Internal Spanning Tree (IST) instance

15. Which MST instance is the IST?

IST is Instance 0. Instances 1 through 15 are available for other use.

16. When an MST region meets a PVST+ domain, how is each MST instance propagated into PVST+?

The BPDUs from each instance are replicated and sent into all the appropriate VLANs in the PVST+ switch.

17. Is it wise to assign VLANs to MST Instance 0? Why or why not?

No. By default, all VLANs are mapped to Instance 0, the IST. You should select the number of instances needed and map all active VLANs to them. Otherwise, you can't have full control over the topologies independent of IST and CST.

18. The commands have just been entered to define an MST region on a switch. You are still at the MST configuration prompt. What command must you enter to commit the MST changes on the switch?

exit (When the MST configuration mode is exited, the changes are committed immediately.)

19. Which of the following methods can you use to assign or propagate MST configuration information to other switches?

    a. Manual configuration
    b. CDP
    c. VTP
    d. MSTP

    Answer: a

20. A switch can interact with both 802.1D and RSTP. Can it run both PVST+ and MST simultaneously?

    No. A switch can run either PVST+ or MST. If a switch is running MST, it can interact and interoperate with PVST+, 802.1D, and RSTP.

Chapter 13

"Do I Know This Already?"

1. d
2. a
3. a
4. b
5. c
6. c
7. c
8. d
9. c
10. c
11. d
12. c